

Real-time Intelligence with Microsoft Fabric

Paul Andrew
Technical Strategist



Cloud Formations

Paul Andrew



Co-Founder & Director
Chief Technology Officer



/mrpaulandrew



@mrpaulandrew



In/mrpaulandrew

- Mentor | Author
- Speaker | Podcast Host
- Event Organiser

SQL Server 2000



Thank you to all our Sponsors!



Microsoft

Unit8™

IT-LOGIX
BUSINESS INTELLIGENCE

BSG GROUP
DATA ANALYTICS




b.telligent

Corporate Software

pmOne



redgate



Real-time Intelligence with Microsoft Fabric

Paul Andrew
Technical Strategist



Cloud Formations

The Microsoft Fabric Icon Game

Synapse Data Engineering My workspace

Search

Fabric Trial: 20 days left

Data Activator
Detect patterns and conditions in your Power BI reports and streaming data, and then take actions such as alert users or kick-off workflows.

- Reflex (Preview)**
Monitor datasets, queries, and event streams for patterns to trigger actions and alerts.
- Reflex sample (Preview)**
Monitor datasets, queries, and event streams for patterns to trigger actions and alerts with sample data.

Data Engineering
Create a lakehouse and operationalize your workflow to build, transform, and share your data estate.

- Lakehouse**
Store big data for cleaning, querying, reporting, and sharing.
- Notebook**
Explore data and build machine learning solutions with Apache Spark applications.
- Environment**
Set up shared libraries, Spark compute settings, and resources for notebooks and Spark job definitions.
- Spark Job Definition**
Define, schedule, and manage your Apache Spark jobs for big data processing.
- API for GraphQL™ (Preview)**
Create an API for GraphQL to easily connect your applications to Fabric data sources.

Data Factory
Empower your organization to get value from data faster than ever.

- Dataflow Gen2**
Prep, clean, and transform data.
- Data pipeline**
Ingest data at scale and schedule data workflows.
- Data workflow (Preview)**
Simplifies the creation and management of Data workflows (powered by Apache Airflow) on which you can operate end-to-end data pipelines at scale.

Data Science
Use machine learning to detect trends, identify outliers, and predict values from your data. [Learn more](#)

- ML model**
Use machine learning models to predict outcomes and detect anomalies in data.
- Experiment**
Create, run, and track development of multiple models for validating hypotheses.
- Notebook**
Explore data and build machine learning solutions with Apache Spark applications.
- Environment**
Set up shared libraries, Spark compute settings, and resources for notebooks and Spark job definitions.

Data Warehouse
Provide strategic insights from multiple sources into your entire business. [Learn more](#)

- Warehouse**
Provide strategic insights from multiple sources into your entire business.
- Mirrored Azure SQL Database ...**
Easily replicate data from an existing source into an analytics-friendly format.
- Mirrored Snowflake (preview)**
Easily replicate data from an existing source into an analytics-friendly format.
- Mirrored Azure Cosmos DB (p...**
Easily replicate data from an existing source into an analytics-friendly format.

Industry Solutions
Use out-of-the-box industry data solutions and resources.

- Sustainability solutions (Pre...**
Unify and process disparate environmental, social, and governance (ESG) data for regulatory disclosures, actionable insights, and analytics.
- Retail solutions (Preview)**
With Microsoft retail data solutions, you can manage retail data at scale to improve customer experience and drive operational efficiency across the organization.
- Healthcare solutions (Preview)**
Use advanced AI analytics to help close care gaps, generate new insights, enhance patient care, and improve outcomes.

Power BI
Use data to find insights, track progress, and make decisions faster. [Learn more](#)

- Report**
Create an interactive presentation of your data.
- Paginated report**
Display tabular data in a report that's easy to print and share.
- Scorecard**
Define, track, and share key metrics for your organization.
- Dashboard**
Build a single-page data story.
- Streaming dataset**
Build visuals from real-time data.

Microsoft Fabric →

- Power BI
- Data Factory
- Data Activator
- Industry Solutions
- Real-Time Intelligence
- Synapse
 - Data Engineering
 - Data Science
 - Data Warehouse



Microsoft Fabric





Power BI







Data Activator





Reflex





Data Warehouse







Real-time Intelligence





KQL Database








Event House





Notebook





KQL Queryset





Reflex





Real-time Hub



The Microsoft Fabric Icon Game



- Home
- Create
- Browse
- OneLake data hub
- Monitor
- Real-Time hub
- Workspaces
- My workspace

Data Activator

Detect patterns and conditions in your Power BI reports and streaming data, and then take actions such as alert users or kick-off workflows.

Reflex (Preview)
Monitor datasets, queries, and event streams for patterns to trigger actions and alerts.

Reflex sample (Preview)
Monitor datasets, queries, and event streams for patterns to trigger actions and alerts with sample data.

Data Engineering

Create a lakehouse and operationalize your workflow to build, transform, and share your data estate.

Lakehouse
Store big data for cleaning, querying, reporting, and sharing.

Notebook
Explore data and build machine learning solutions with Apache Spark applications.

Environment
Set up shared libraries, Spark compute settings, and resources for notebooks and Spark job definitions.

Spark Job Definition
Define, schedule, and manage your Apache Spark jobs for big data processing.

API for GraphQL™ (Preview)
Create an API for GraphQL to easily connect your applications to Fabric data sources.

Data Factory

Empower your organization to get value from data faster than ever.

Dataflow Gen2
Prep, clean, and transform data.

Data pipeline
Ingest data at scale and schedule data workflows.

Data workflow (Preview)
Simplifies the creation and management of Data workflows (powered by Apache Airflow) on which you can operate end-to-end data pipelines at scale.

Data Science

Use machine learning to detect trends, identify outliers, and predict values from your data. [Learn more](#)

ML model
Use machine learning models to predict outcomes and detect anomalies in data.

Experiment
Create, run, and track development of multiple models for validating hypotheses.

Notebook
Explore data and build machine learning solutions with Apache Spark applications.

Environment
Set up shared libraries, Spark compute settings, and resources for notebooks and Spark job definitions.

Data Warehouse

Provide strategic insights from multiple sources into your entire business. [Learn more](#)

Warehouse
Provide strategic insights from multiple sources into your entire business.

Mirrored Azure SQL Database ...
Easily replicate data from an existing source into an analytics-friendly format.

Mirrored Snowflake (preview)
Easily replicate data from an existing source into an analytics-friendly format.

Mirrored Azure Cosmos DB (p...
Easily replicate data from an existing source into an analytics-friendly format.

Industry Solutions

Use out-of-the-box industry data solutions and resources.

Sustainability solutions (Previe...
Unify and prepare disparate environmental, social, and governance (ESG) data for regulatory disclosures, actionable insights, and analytics.

Retail solutions (Preview)
With Microsoft Retail data solutions, you can manage retail data at scale to improve customer experience and drive operational efficiency across the organization.

Healthcare solutions (Preview)
Use advanced AI analytics to help close care gaps, generate new insights, enhance patient care, and improve outcomes.

Power BI

Use data to find insights, track progress, and make decisions faster. [Learn more](#)

Report
Create an interactive presentation of your data.

Paginated report
Display tabular data in a report that's easy to print and share.

Scorecard
Define, track, and share key metrics for your organization.

Dashboard
Build a single-page data story.

Streaming dataset
Build visuals from real-time data.

Microsoft Fabric →

- Power BI
- Data Factory
- Data Activator
- Industry Solutions
- Real-Time Intelligence
- Synapse
 - Data Engineering
 - Data Science
 - Data Warehouse



Add a slide about the advantages of combining real-time data with batch data in a lambda architecture.

I completed some of your request, but I'll need more practice before I can do everything in it. What else can I help you with?




By using a lambda architecture, we can take advantage of both real-time and batch data to gain a 360-degree view of our data. Real-time data allows us to monitor system performance and detect issues as they arise, while batch data provides a comprehensive view of our data. By combining these two types of data in a lambda architecture, we can gain a 360-degree view of our data.

f
e
a
a
e
S

Real-time Intelligence with Microsoft abric

What | Why | How

Cloud Formations

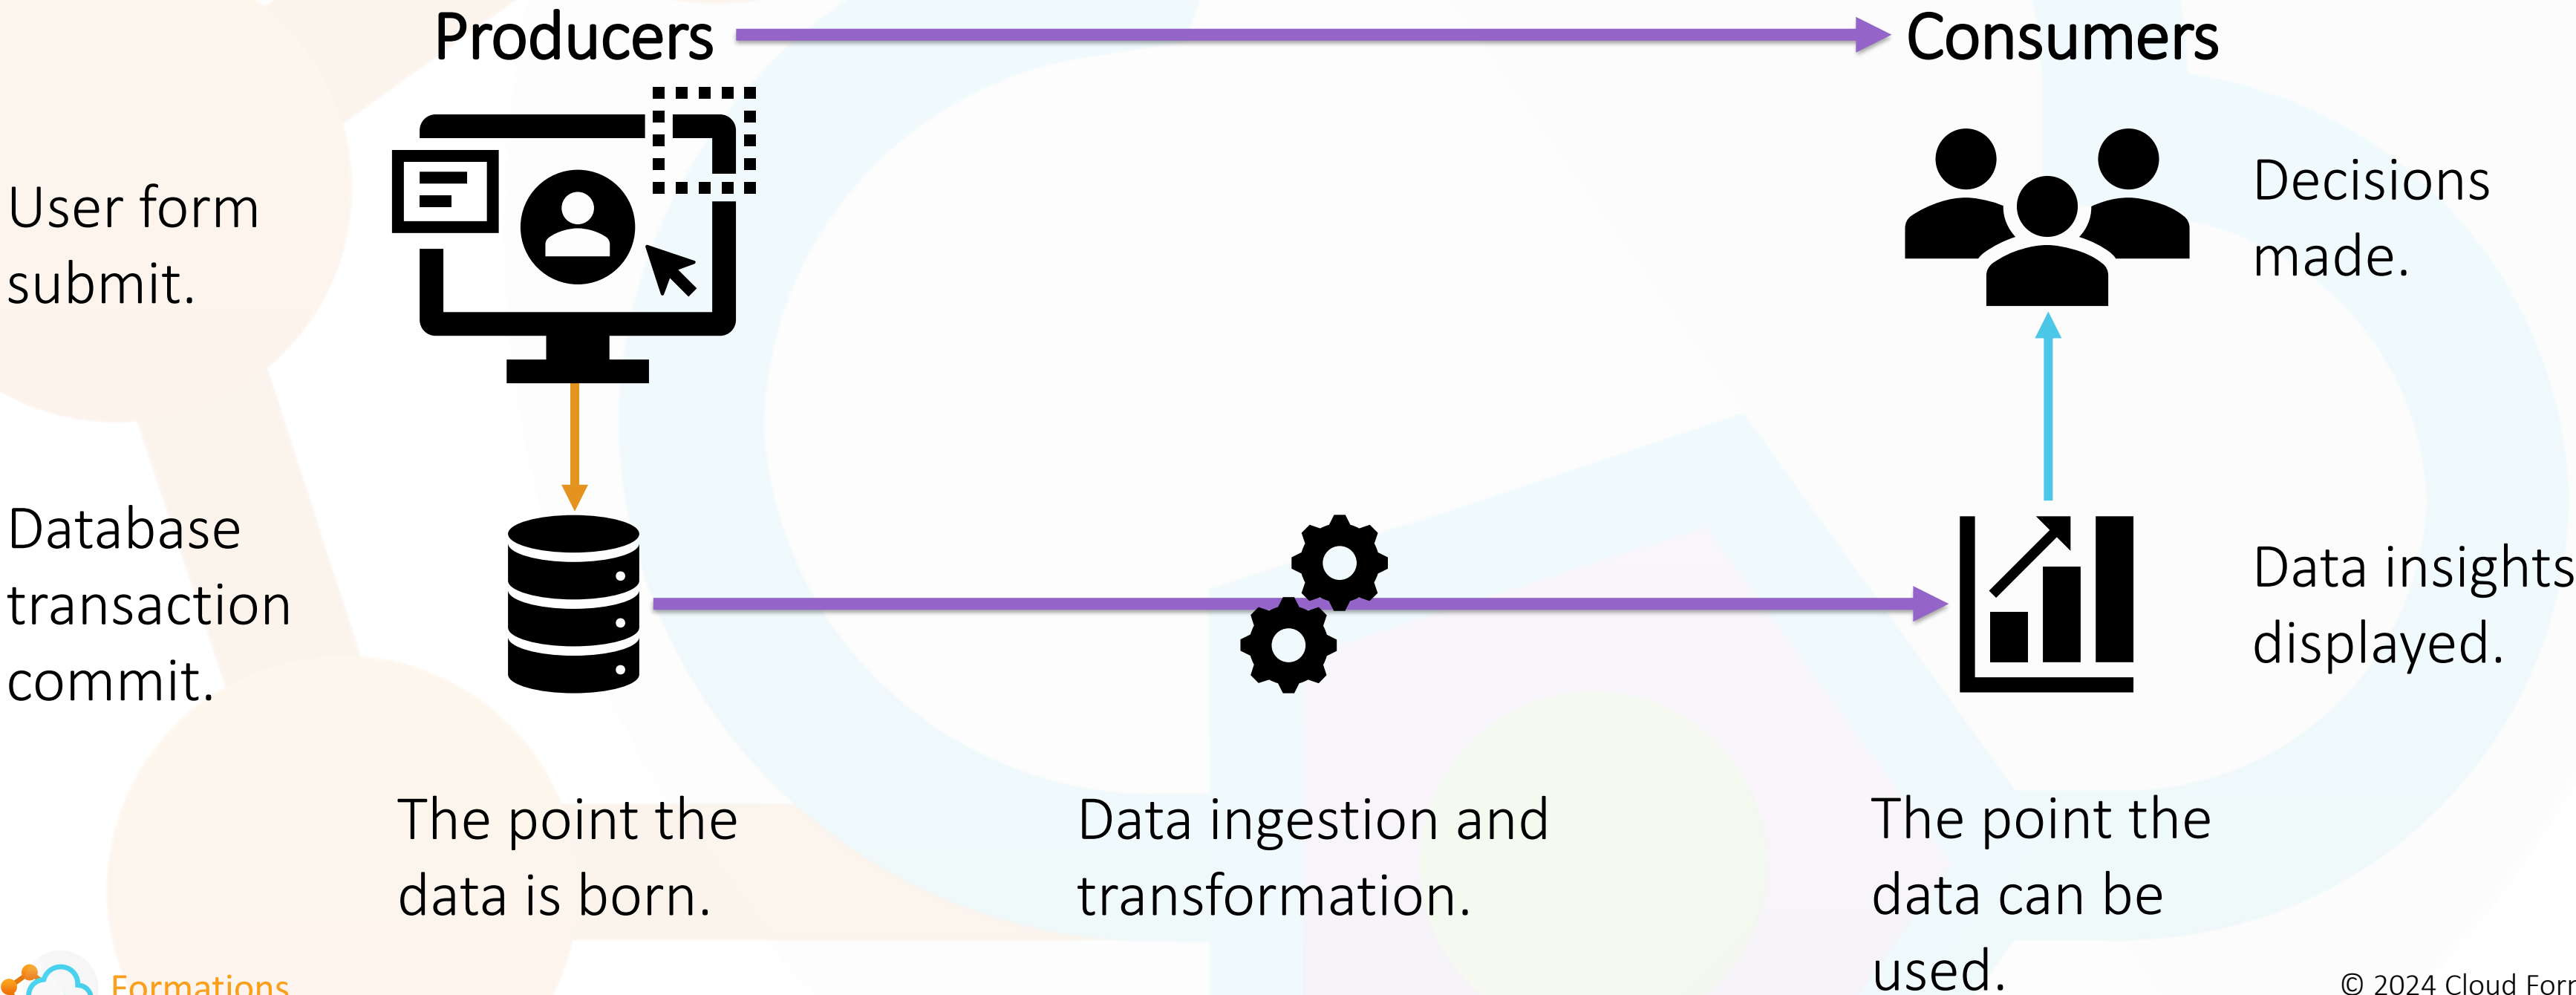


What do we mean by real-time data?

Answer (big data):

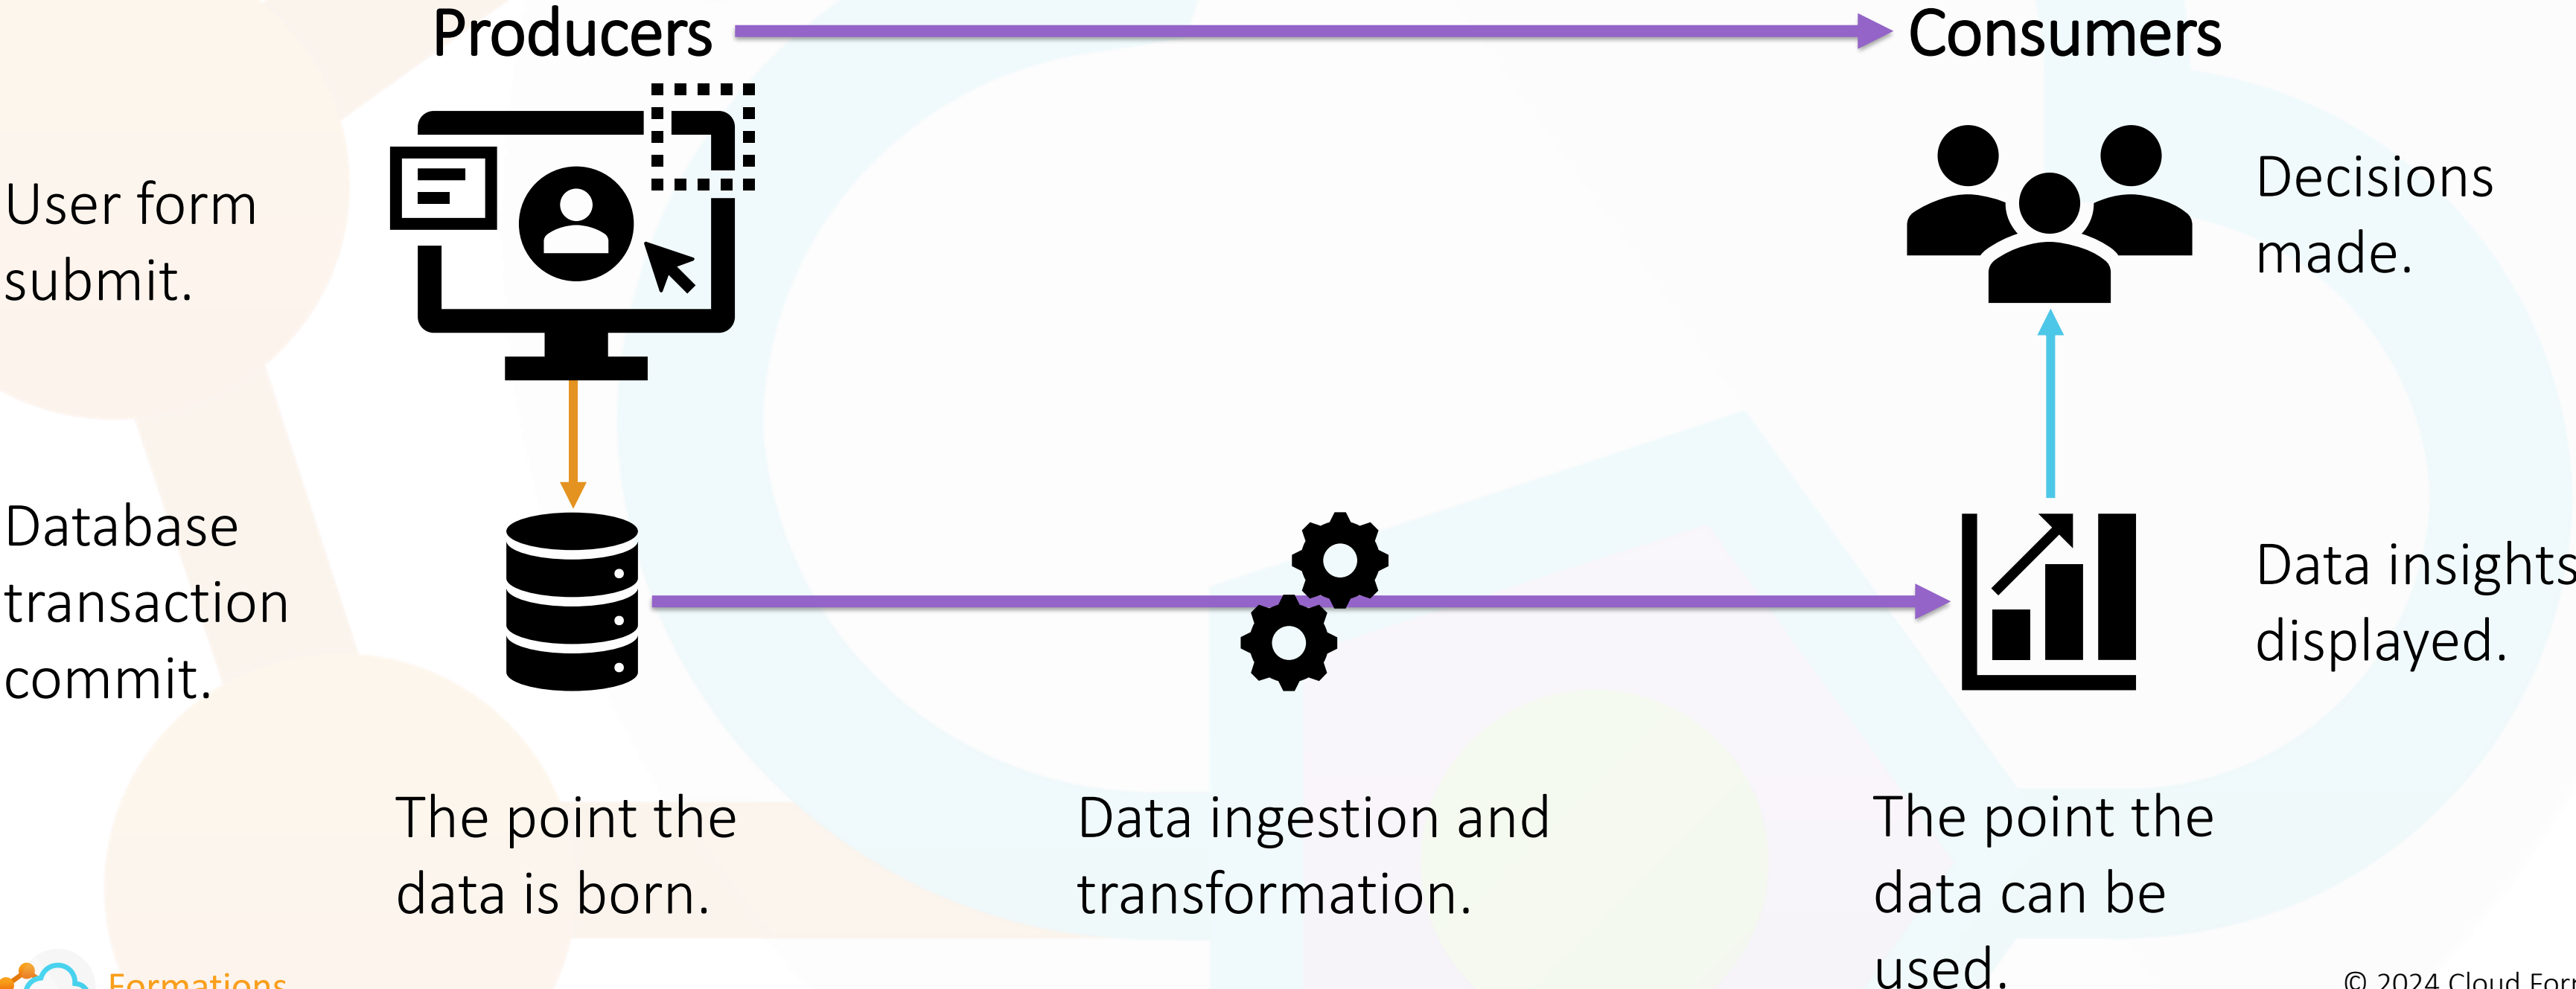
~~“Any data that you cannot process in the time that you have/want using the technology you have.”~~

Buck Woody



What do we mean by real-time data?

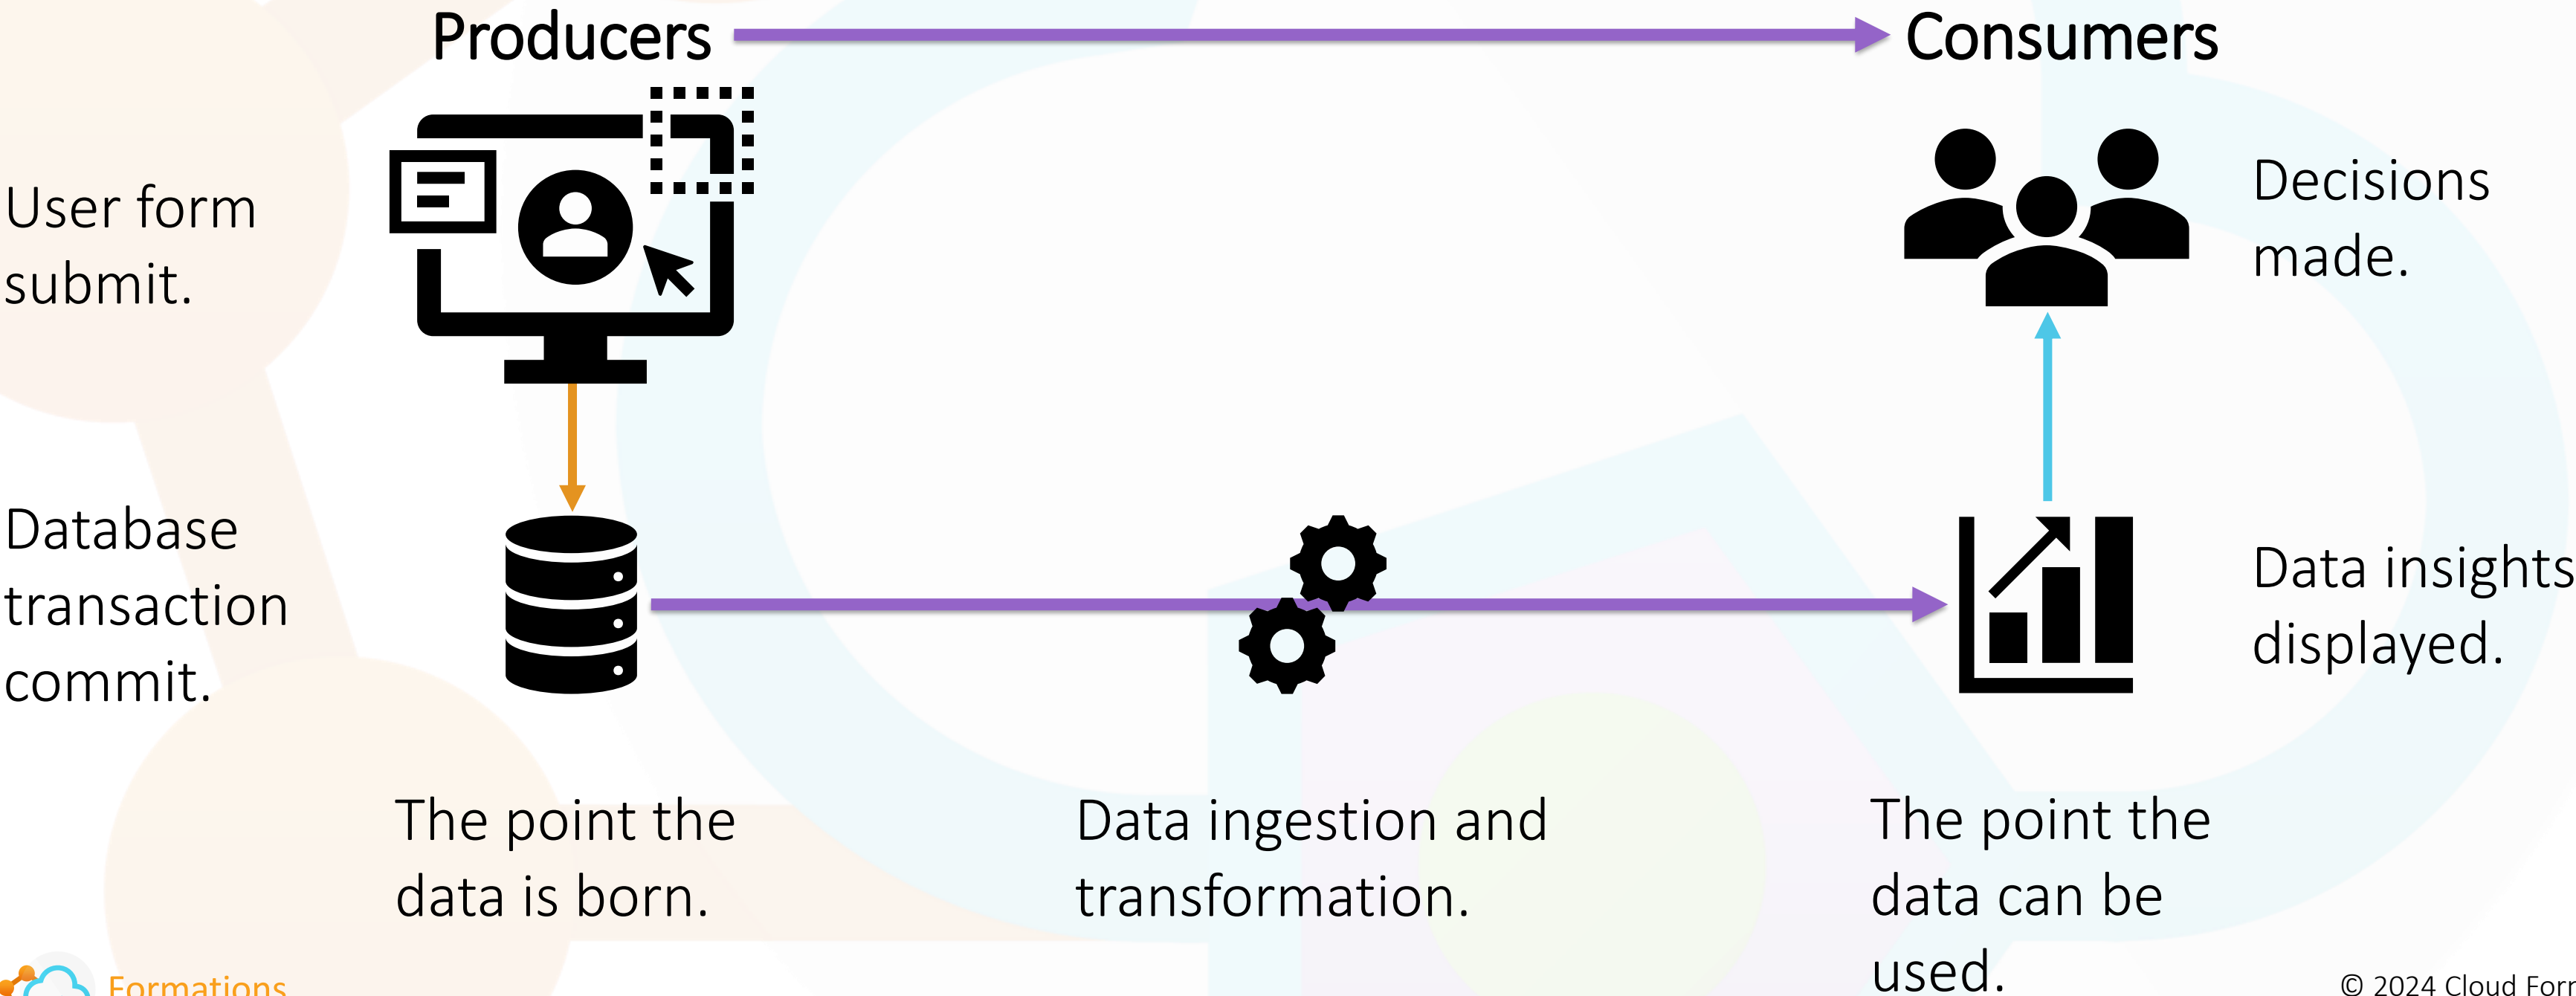
Answer:
“Delivering data from the producer to consumer as fast as possible using the technology you have.”



What do we mean by near real-time data?

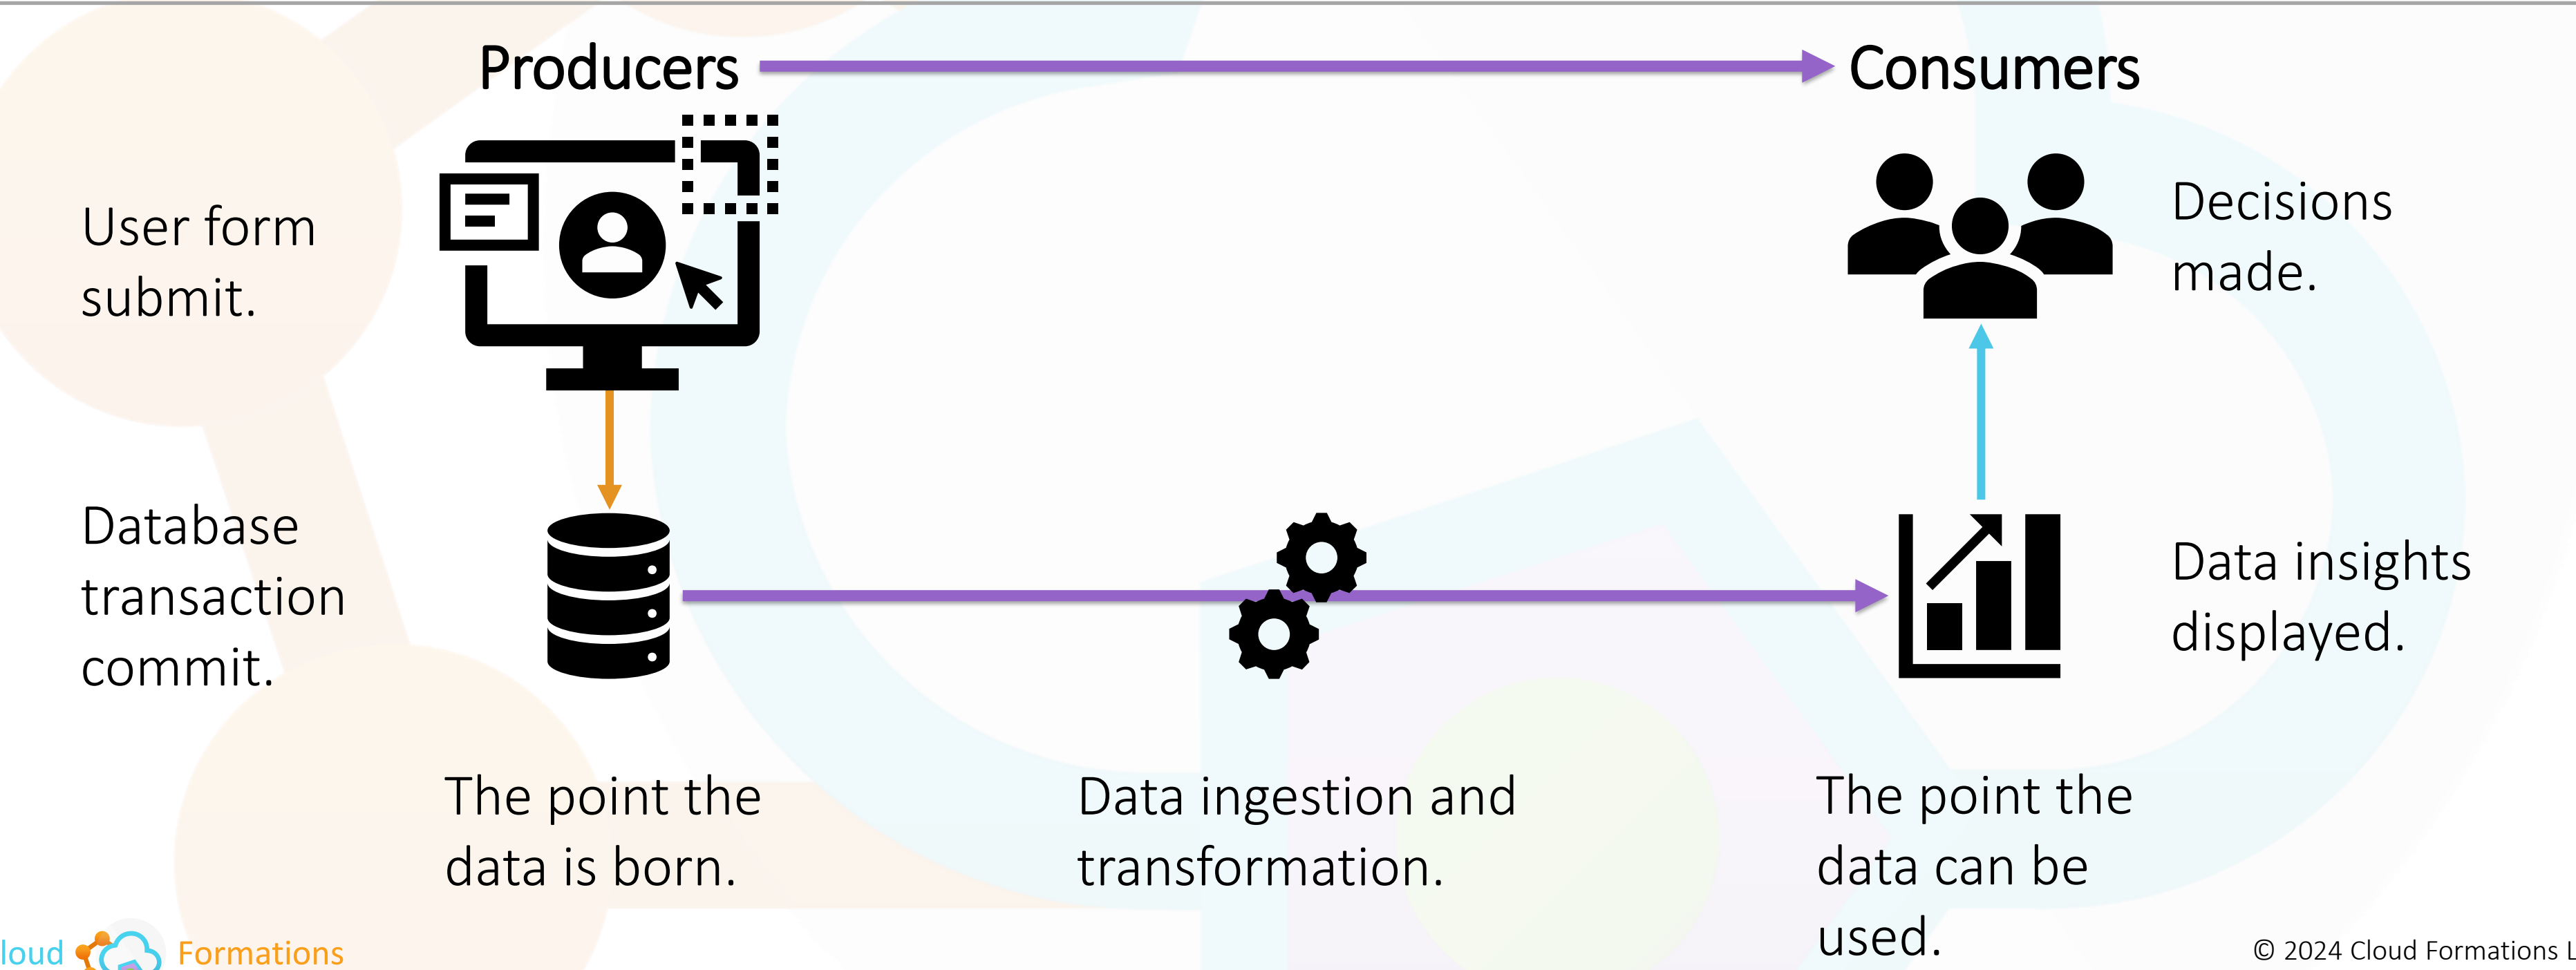
Answer:

“Delivering data from the producer to consumer within 1 minute of it being created (born).”



What do we mean by a data stream?

Answer:
“Data that is constantly flowing from producer to consumer in near real-time.”





Real-Time Intelligence in Microsoft Fabric



Enterprise real-time data platforms

Azure Event Hubs

Azure Event Grid

Azure Stream Analytics

Azure Data Explorer




Self-service reporting and activation

OneLake

Data Activator

Power BI



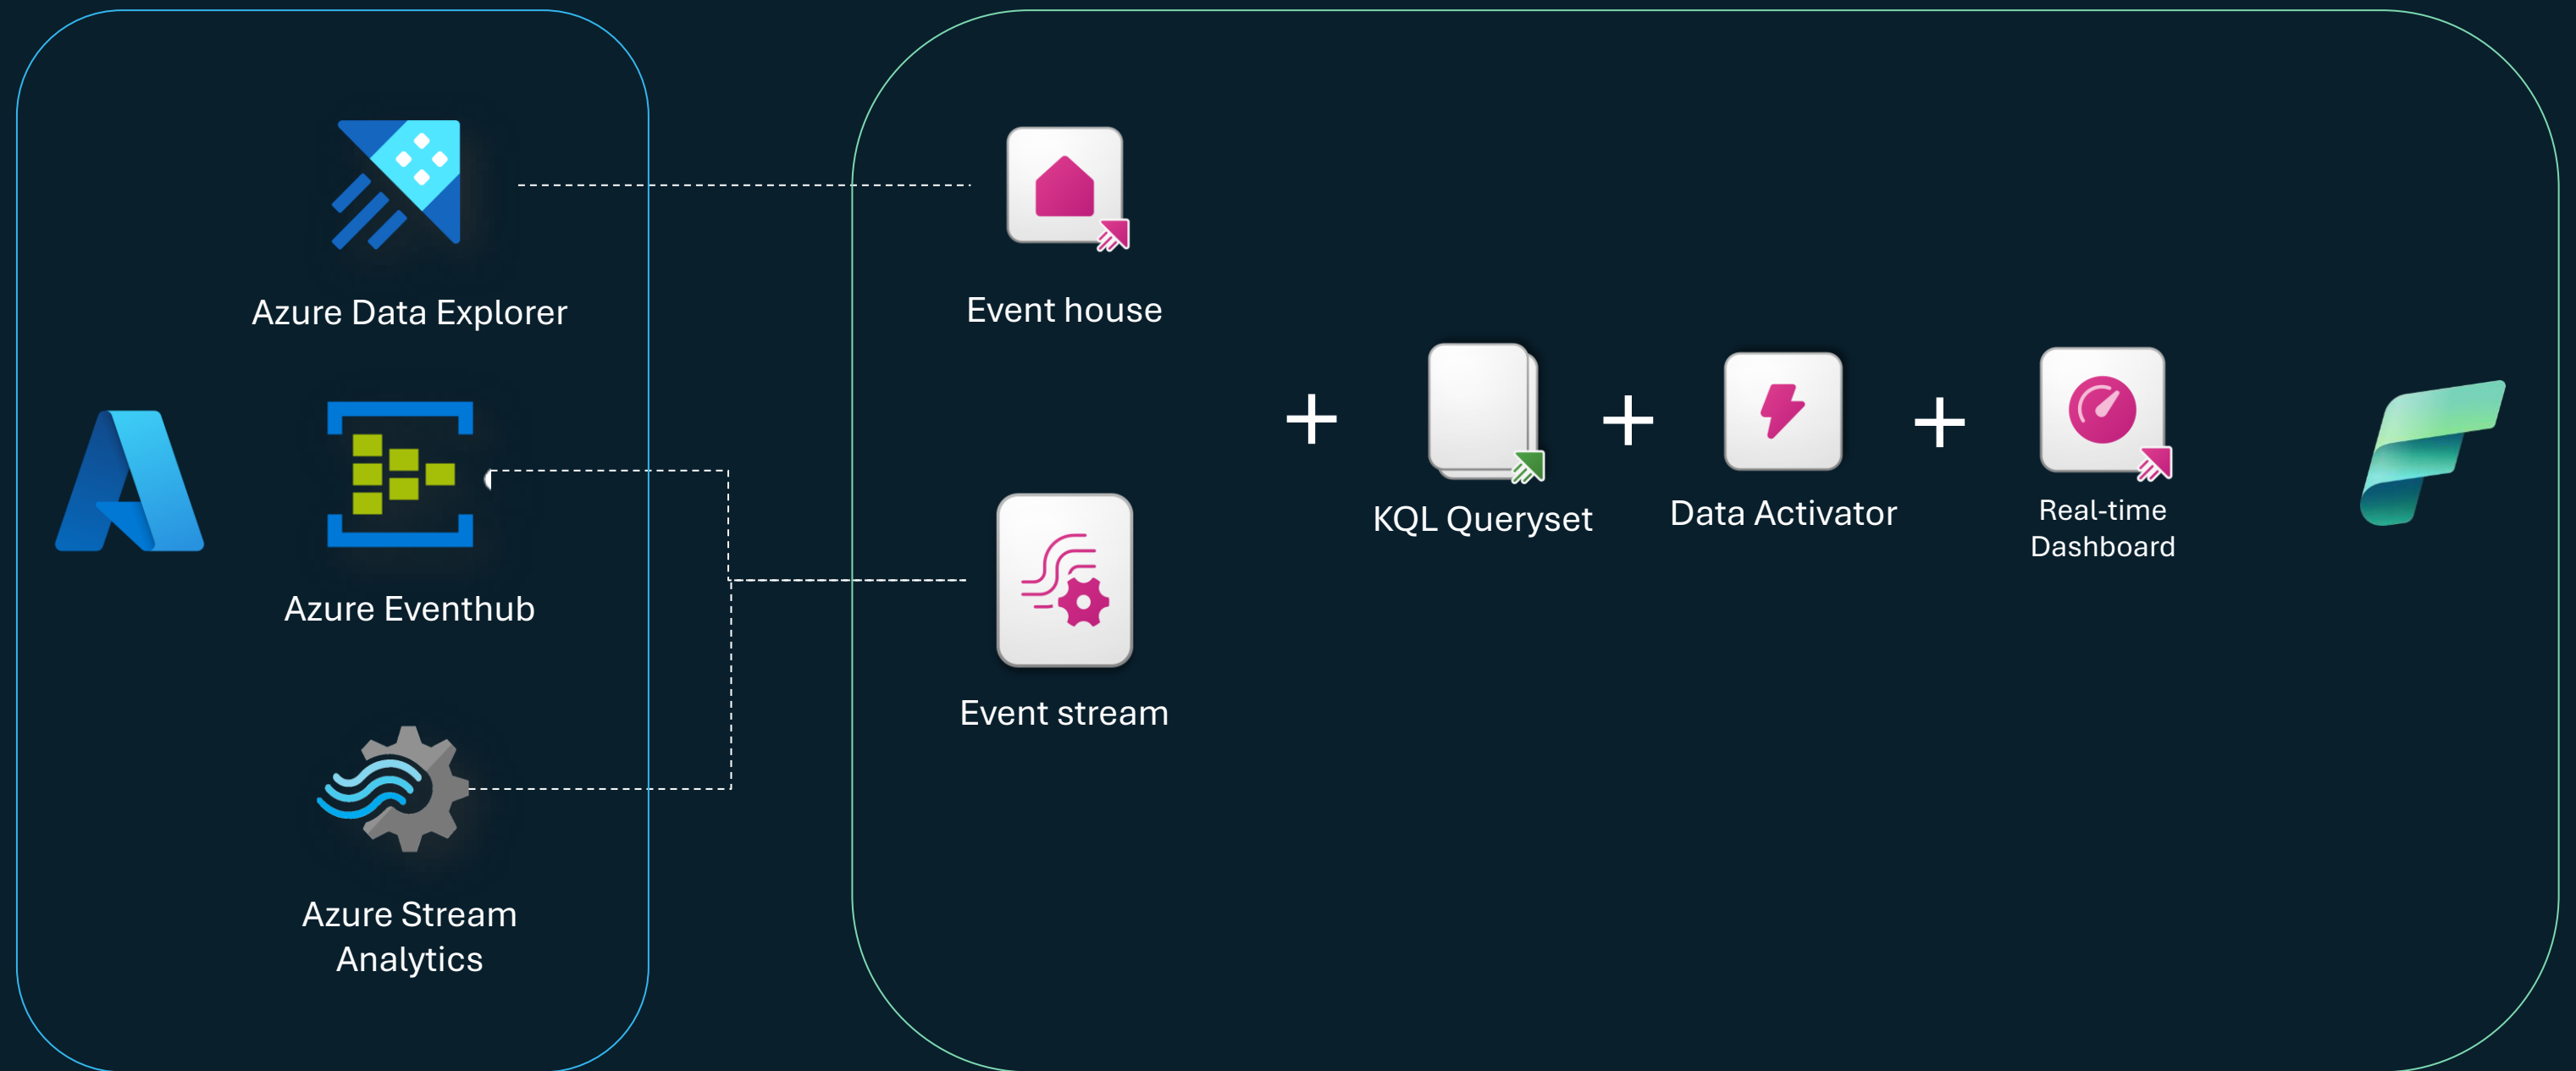
Real-Time Intelligence in Microsoft Fabric

Fully integrated, no/low-code real-time SaaS data platform



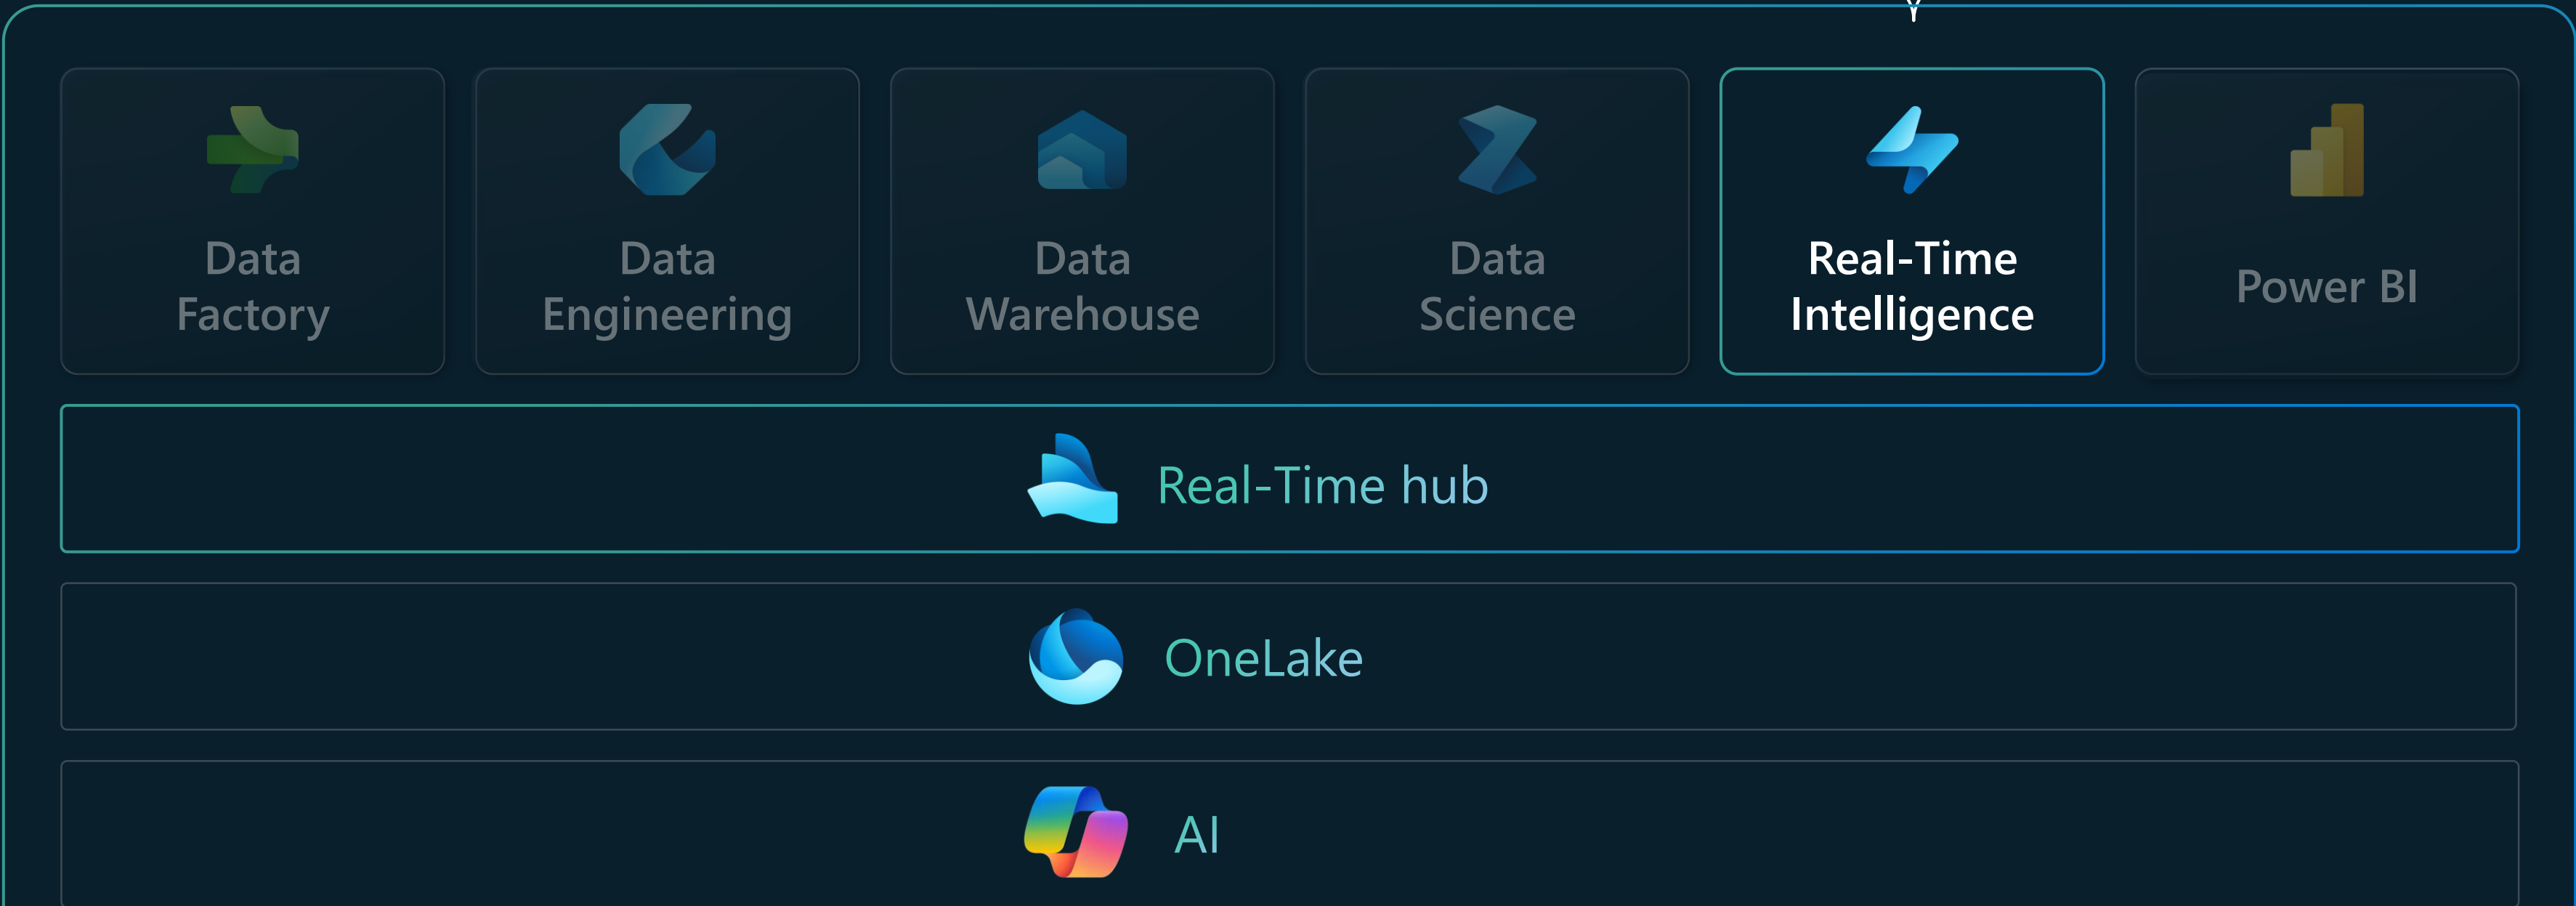
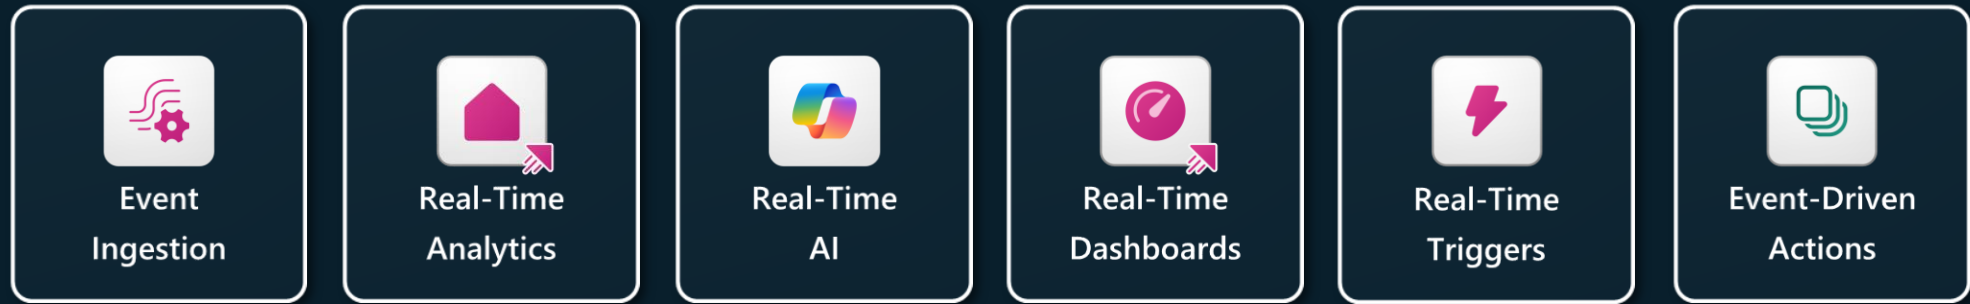
Microsoft Fabric Real-time Intelligence

Evolving Azure services into Microsoft Fabric

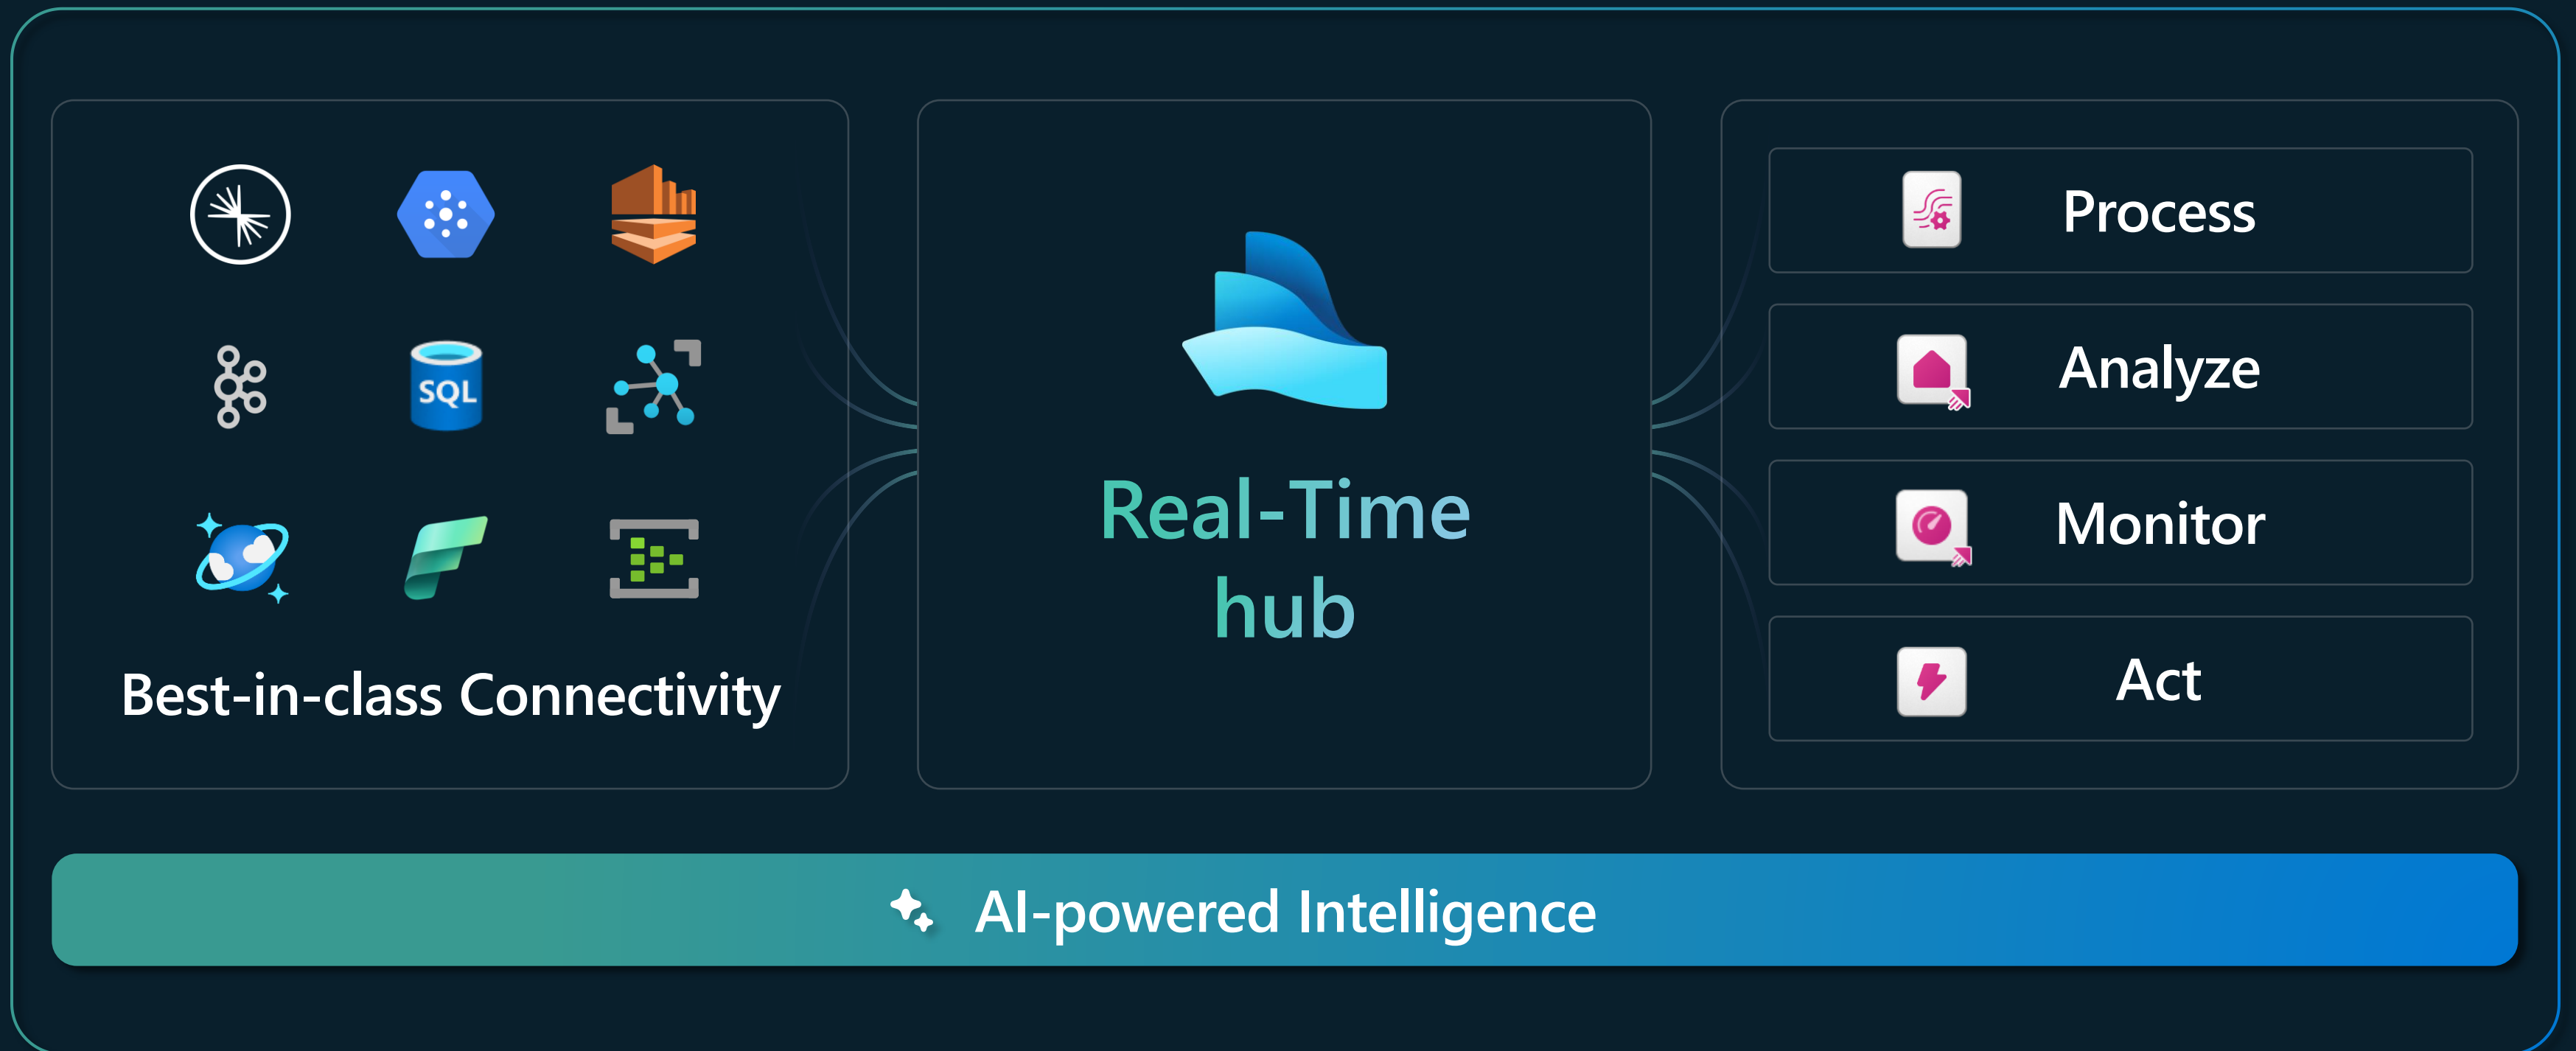




Real-Time Intelligence in Microsoft Fabric




Real-Time Intelligence in Microsoft Fabric





Real-time Intelligence with Microsoft abric

What | Why | **How**

Cloud Formations



Our Use Case



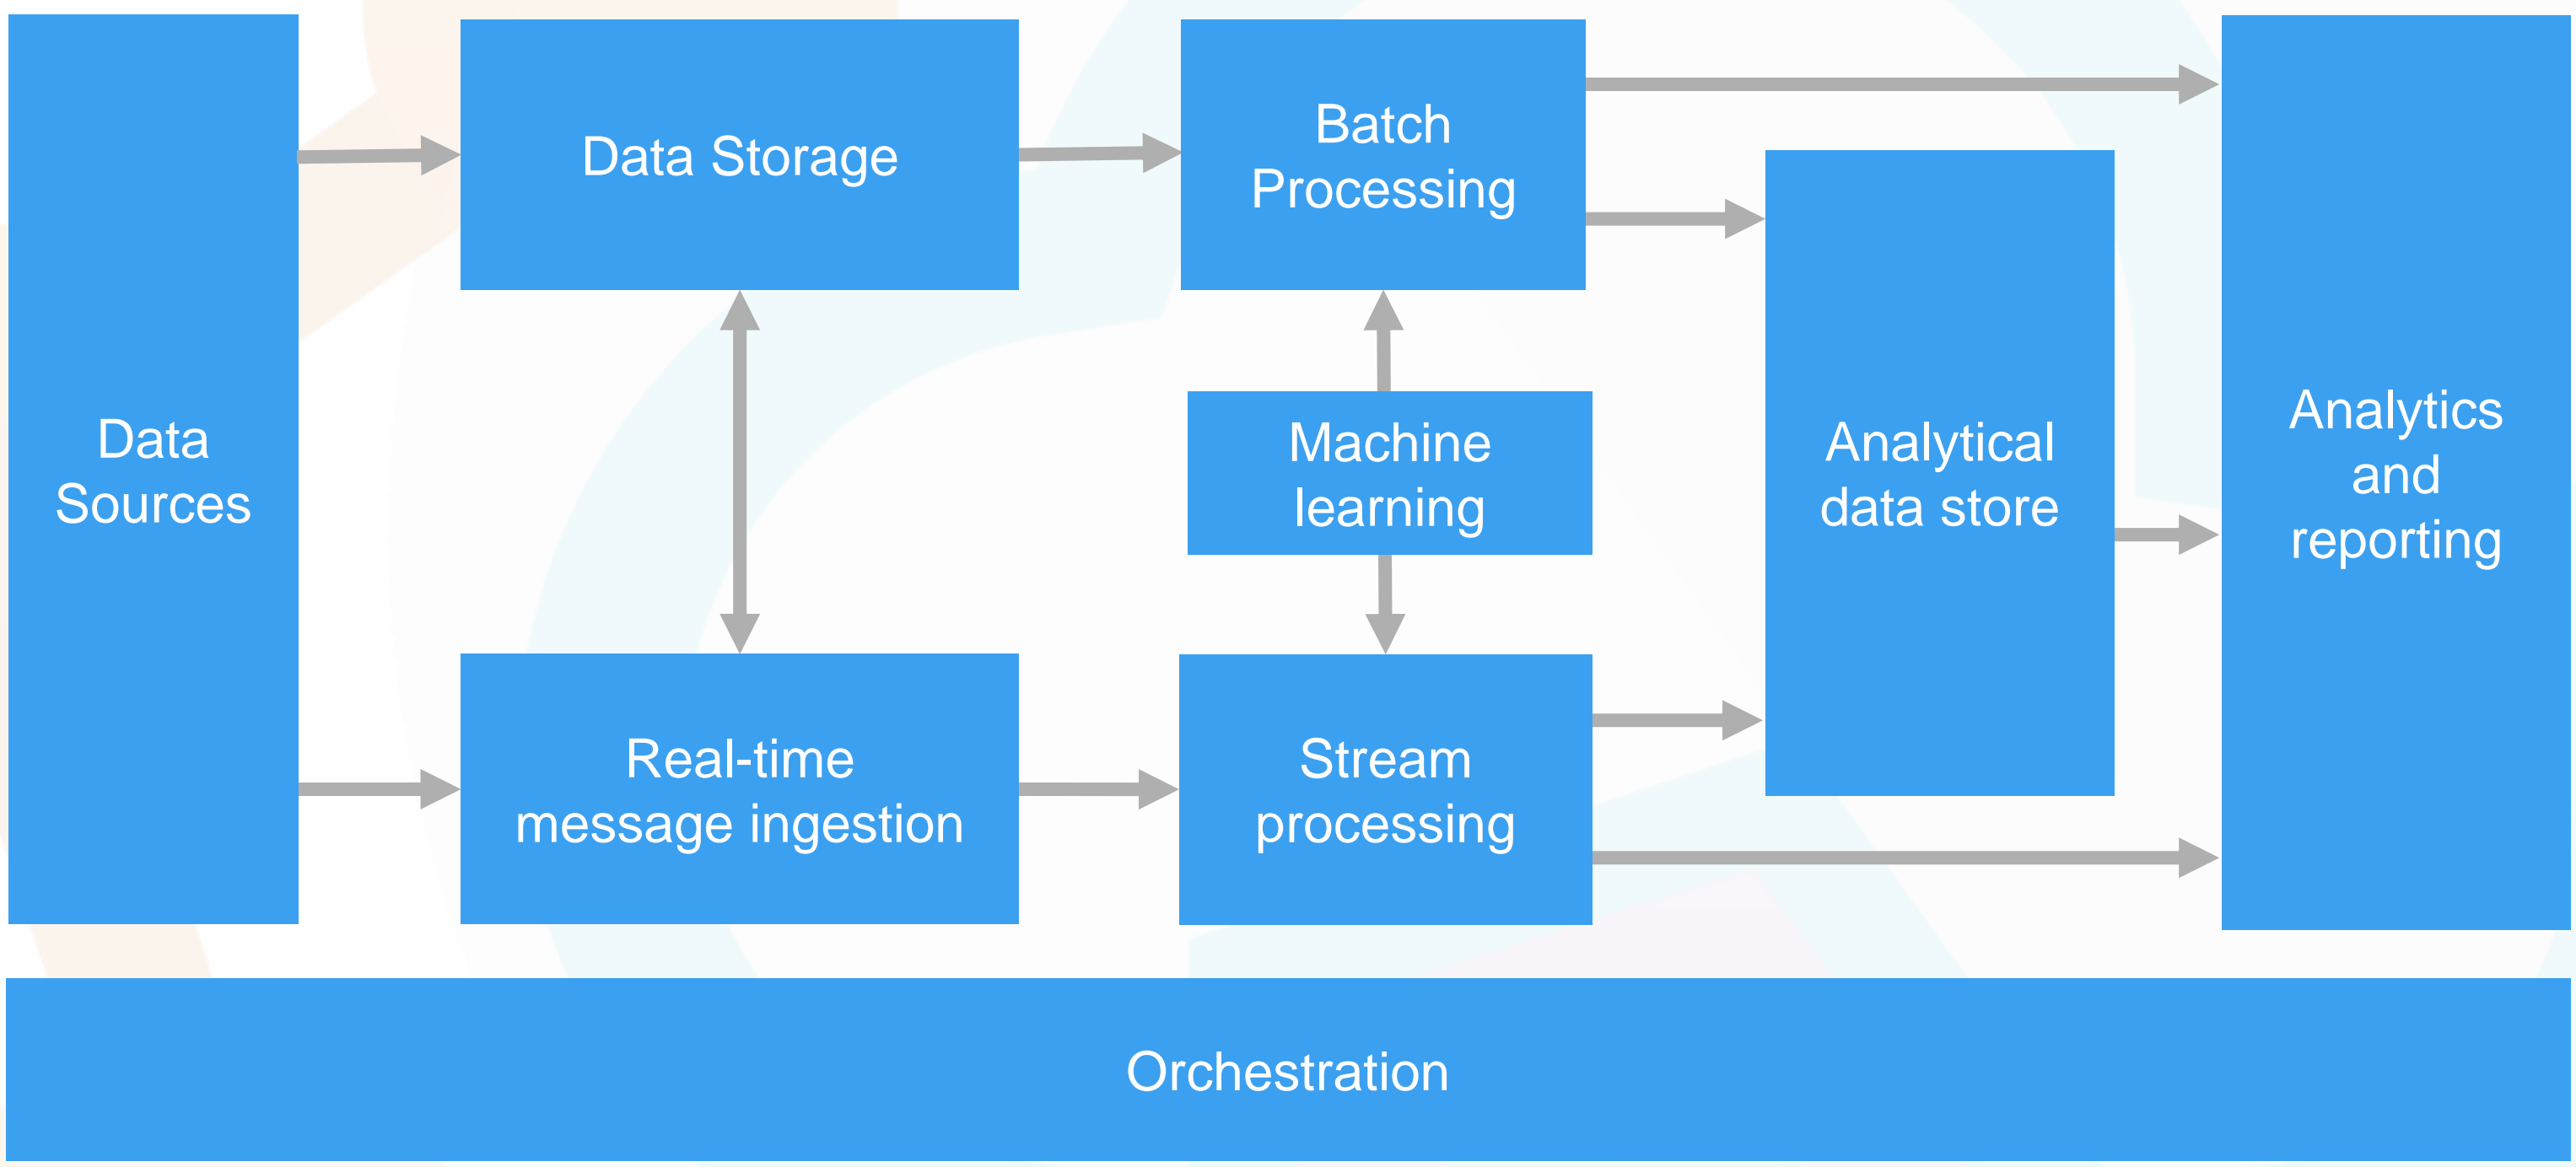
A bike shop front with large glass windows showing bike related products on the shelves inside. The shop front has a sign with the name Adventure Works.

A bike shop called Adventure Works, standing inside looking at the cashier which includes a modern point of sale till system. On the wall behind the cashier are bike related products.

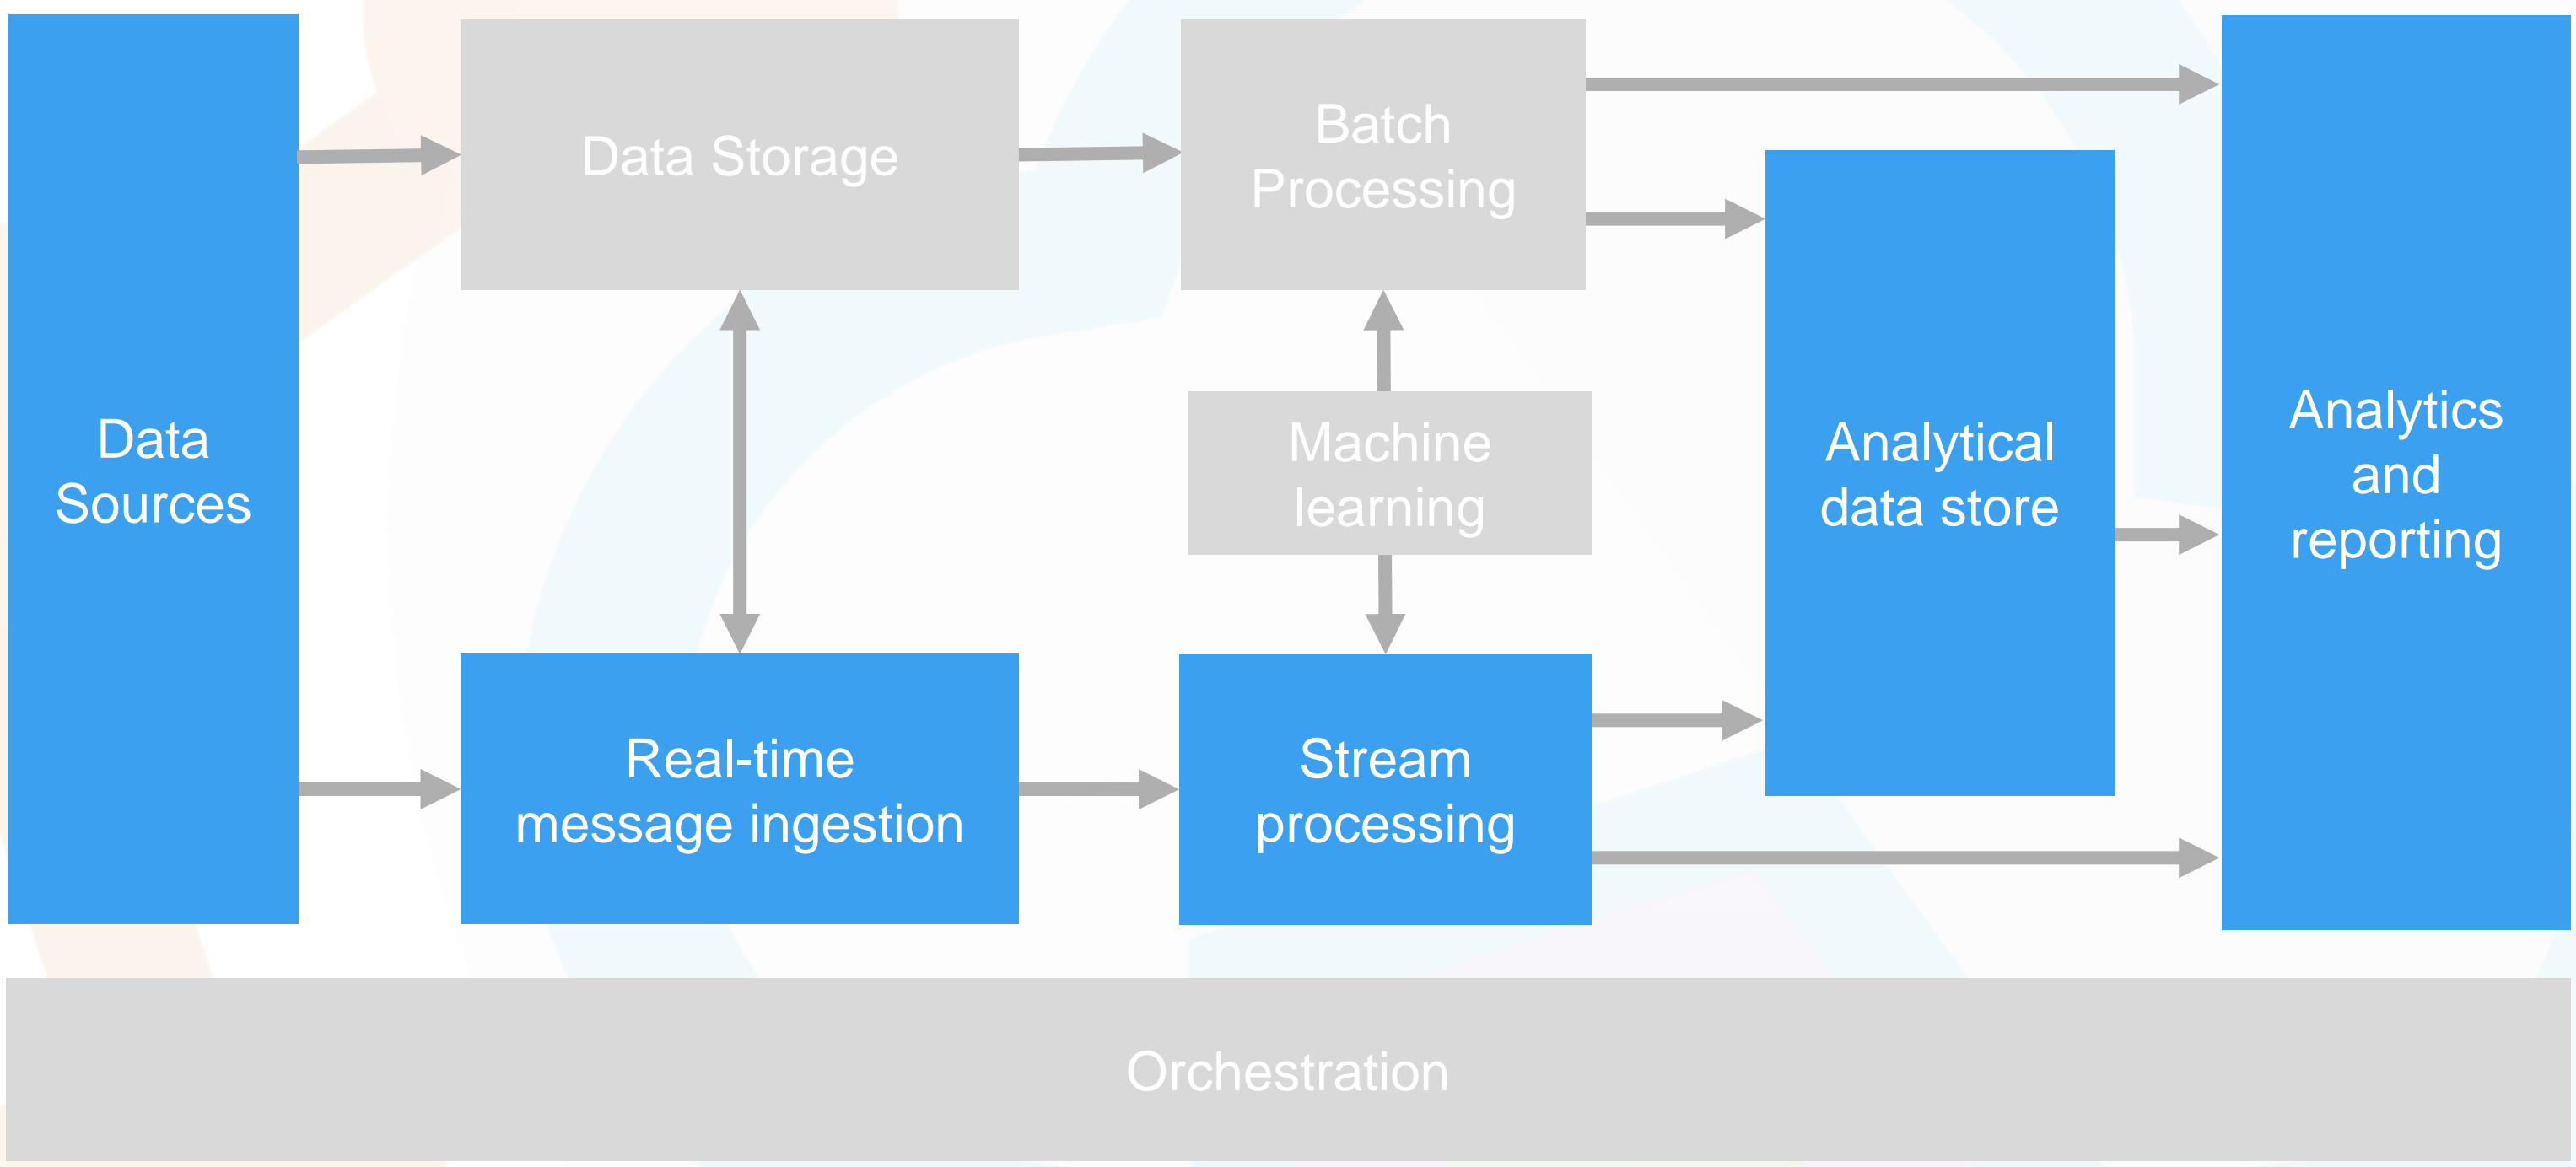
At the corporate head offices of Adventure Works standing in the main operations room. On the wall are large televisions showing a range of data analytics dashboards with charts and technical information.

Inside the office of the CEO at the company Adventure Works, on the desk is a business plan to role out more retail stores across the country based on targeted growth on an analytics dashboard visible on a large TV.

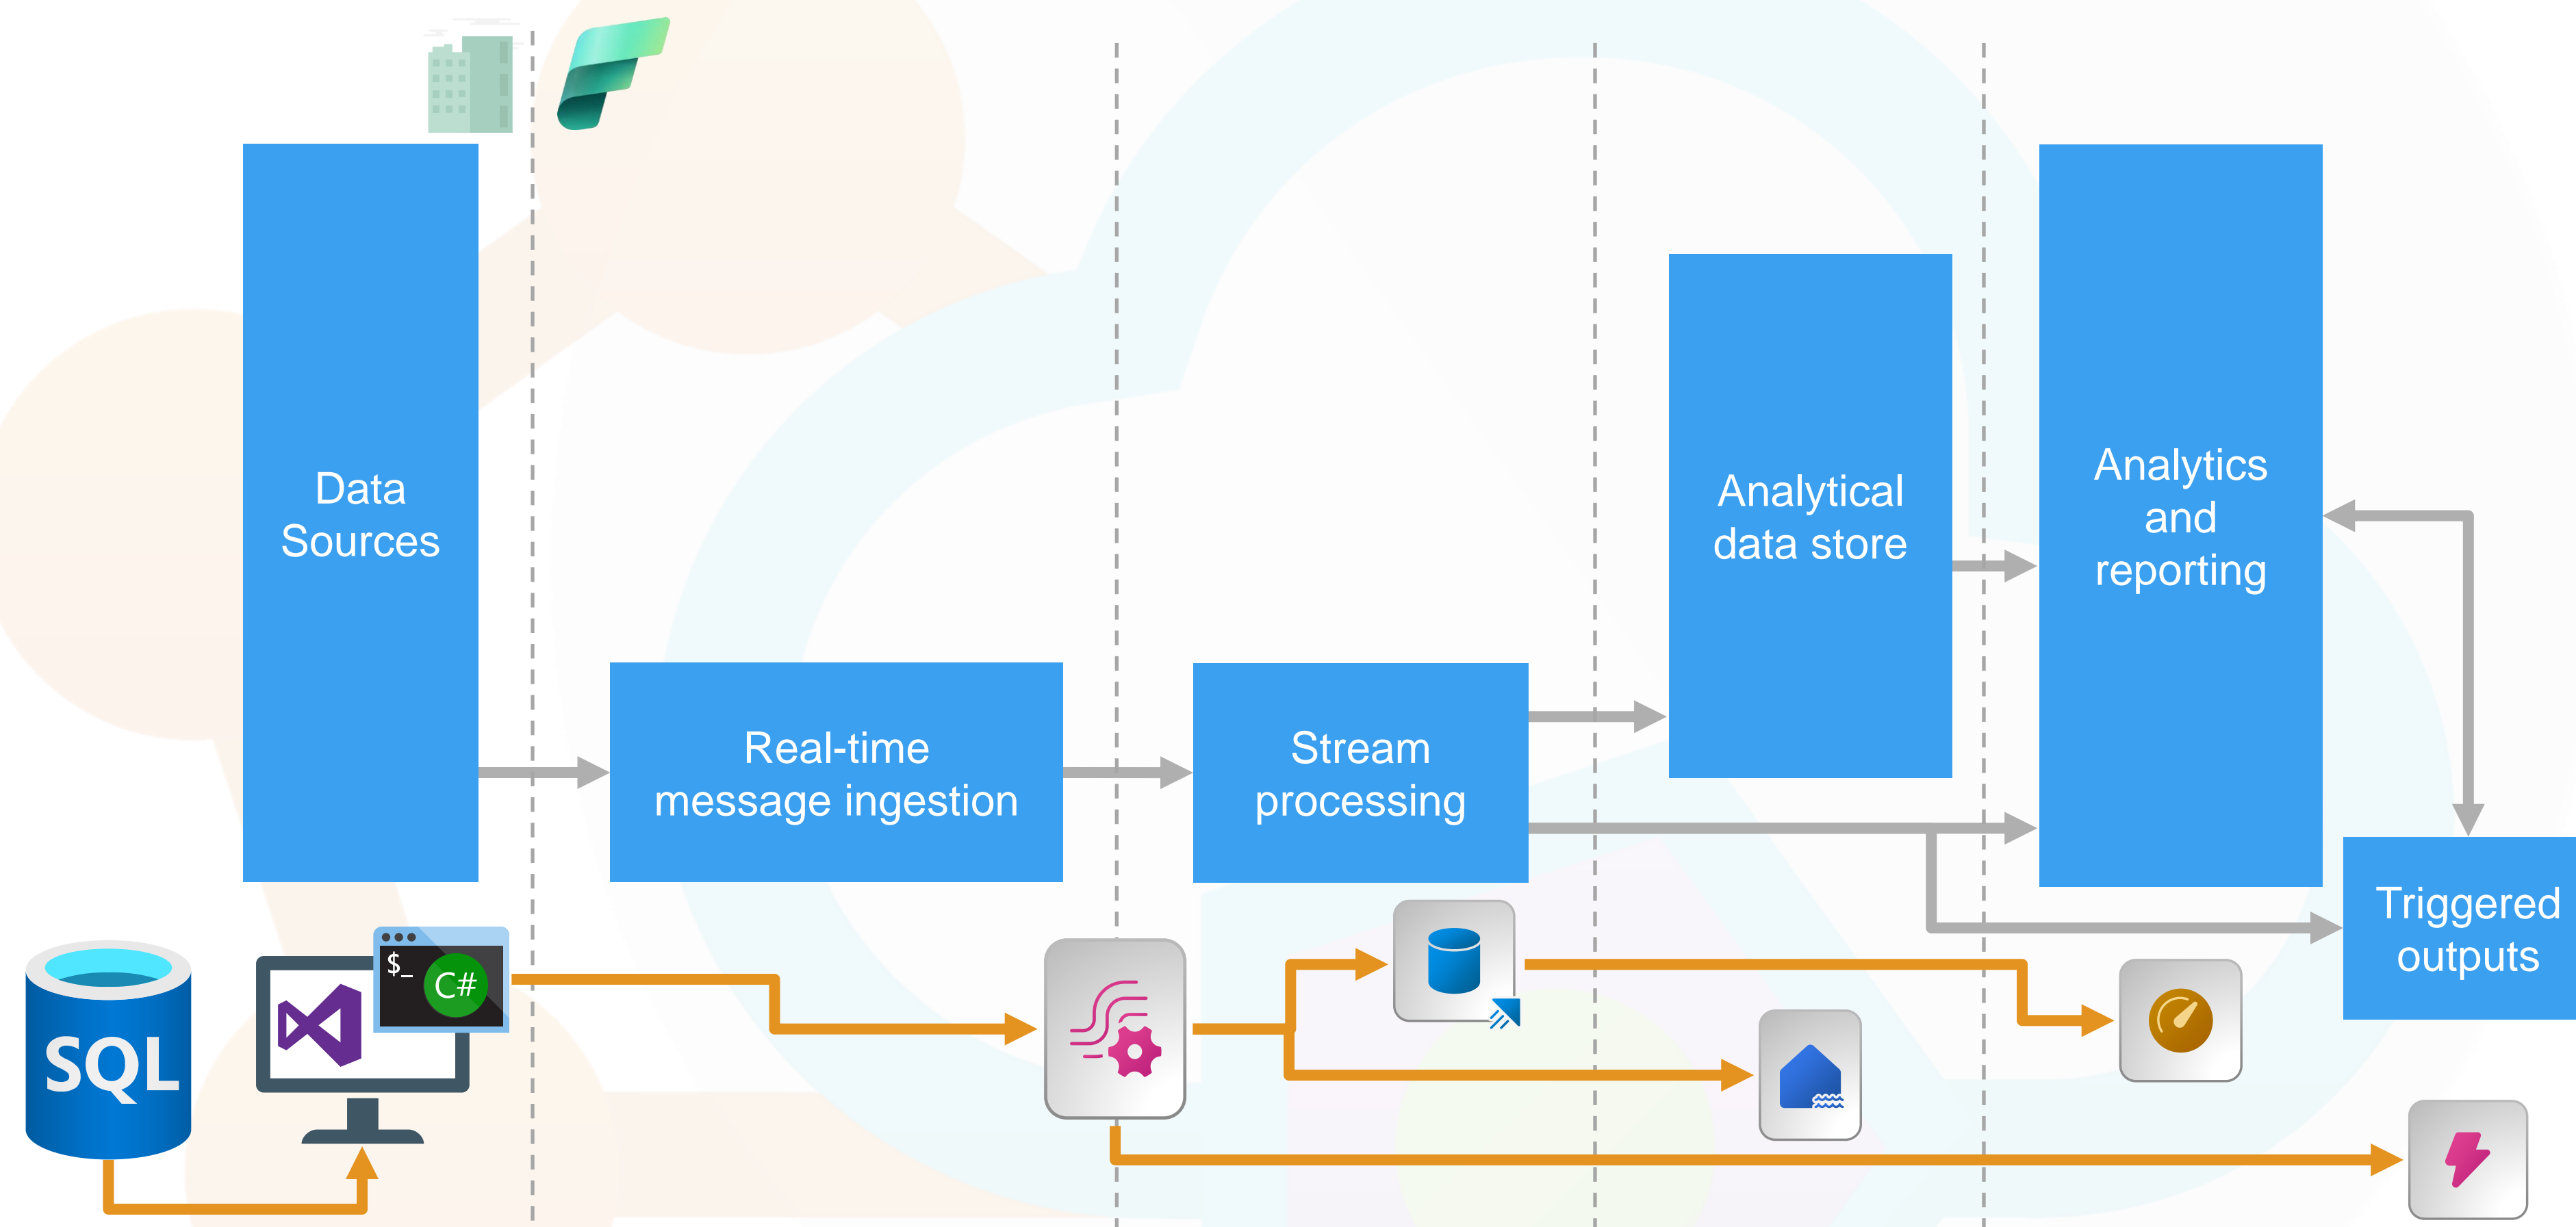
Components of a Big Data Architecture



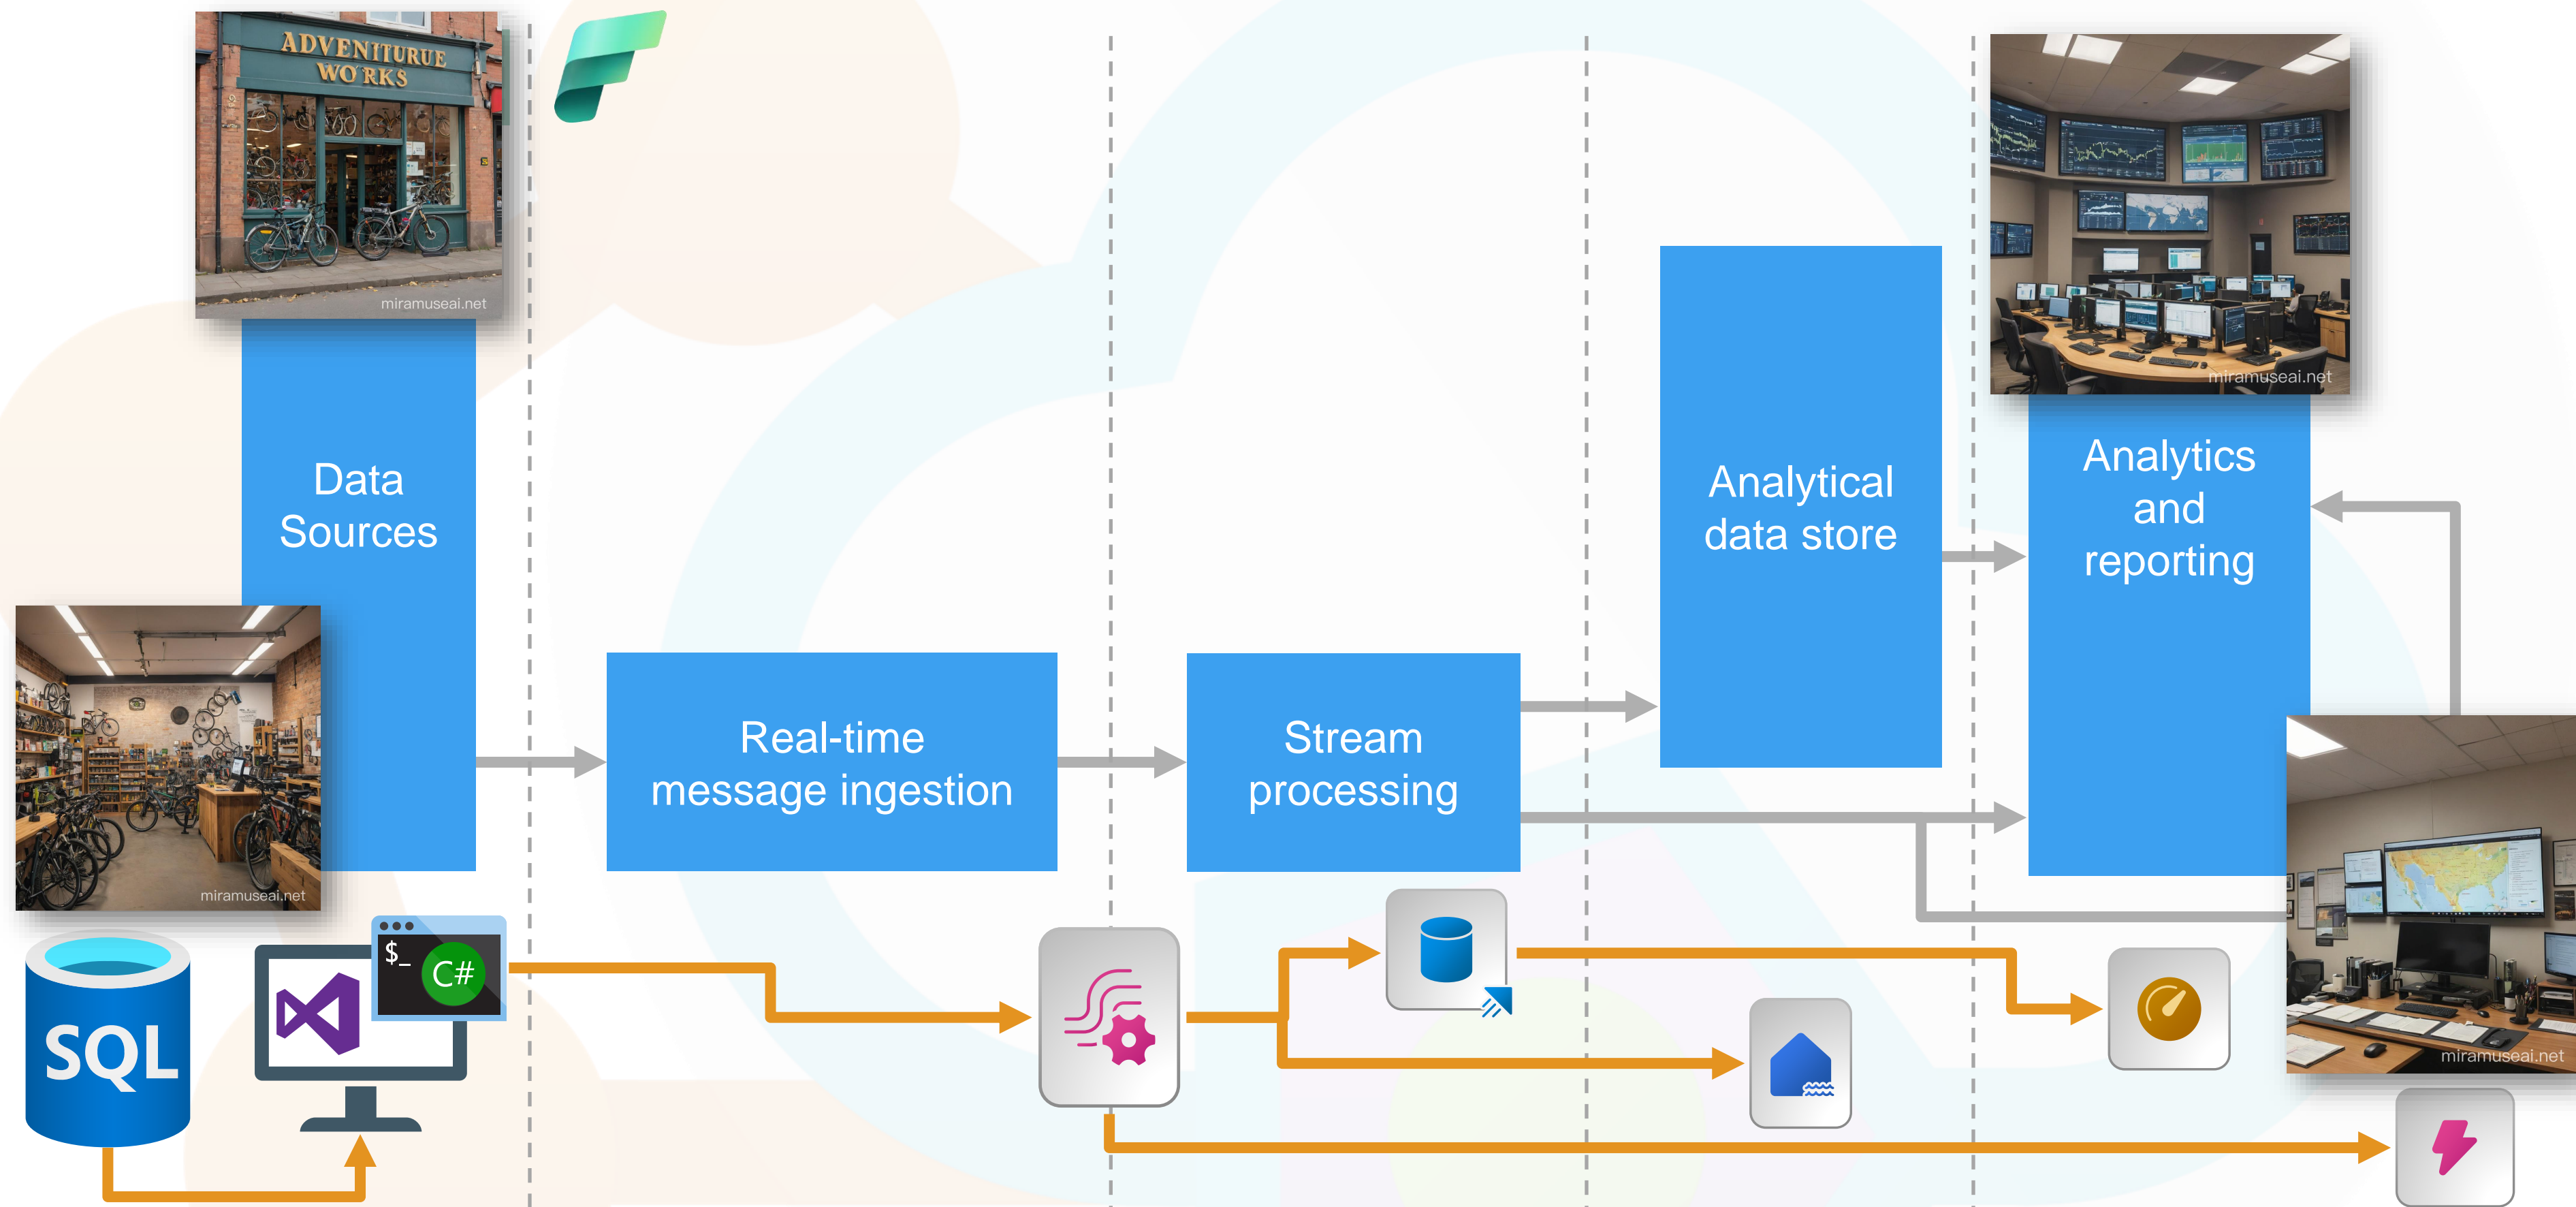
Components of a Big Data Architecture



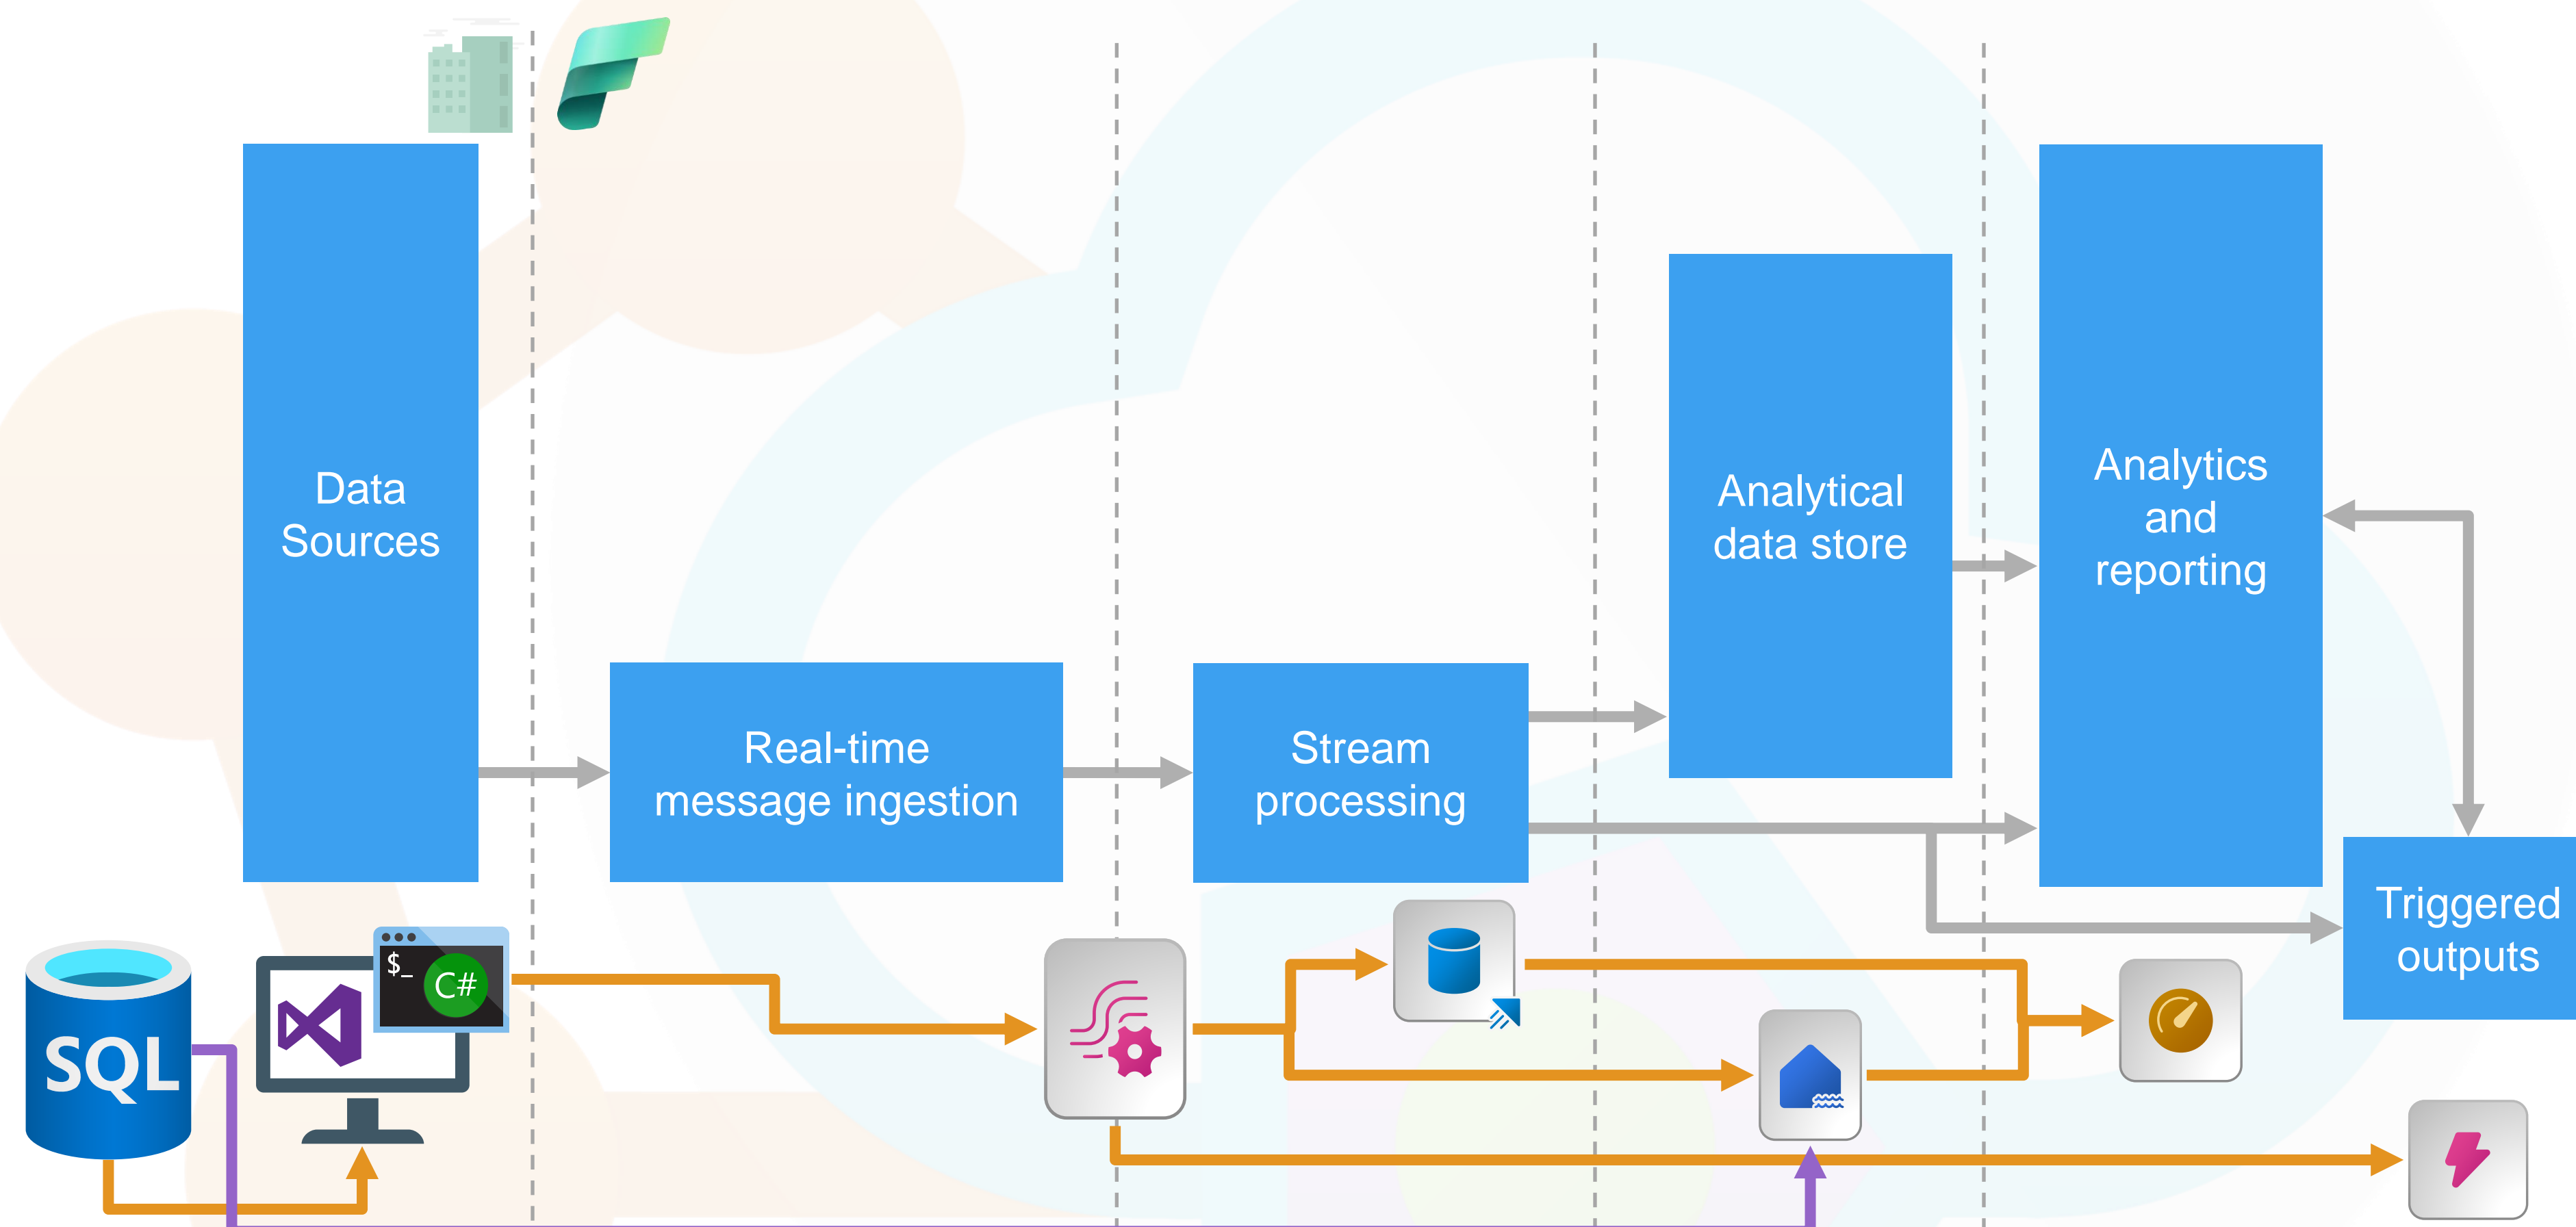
Fabric Real-Time Data Handling



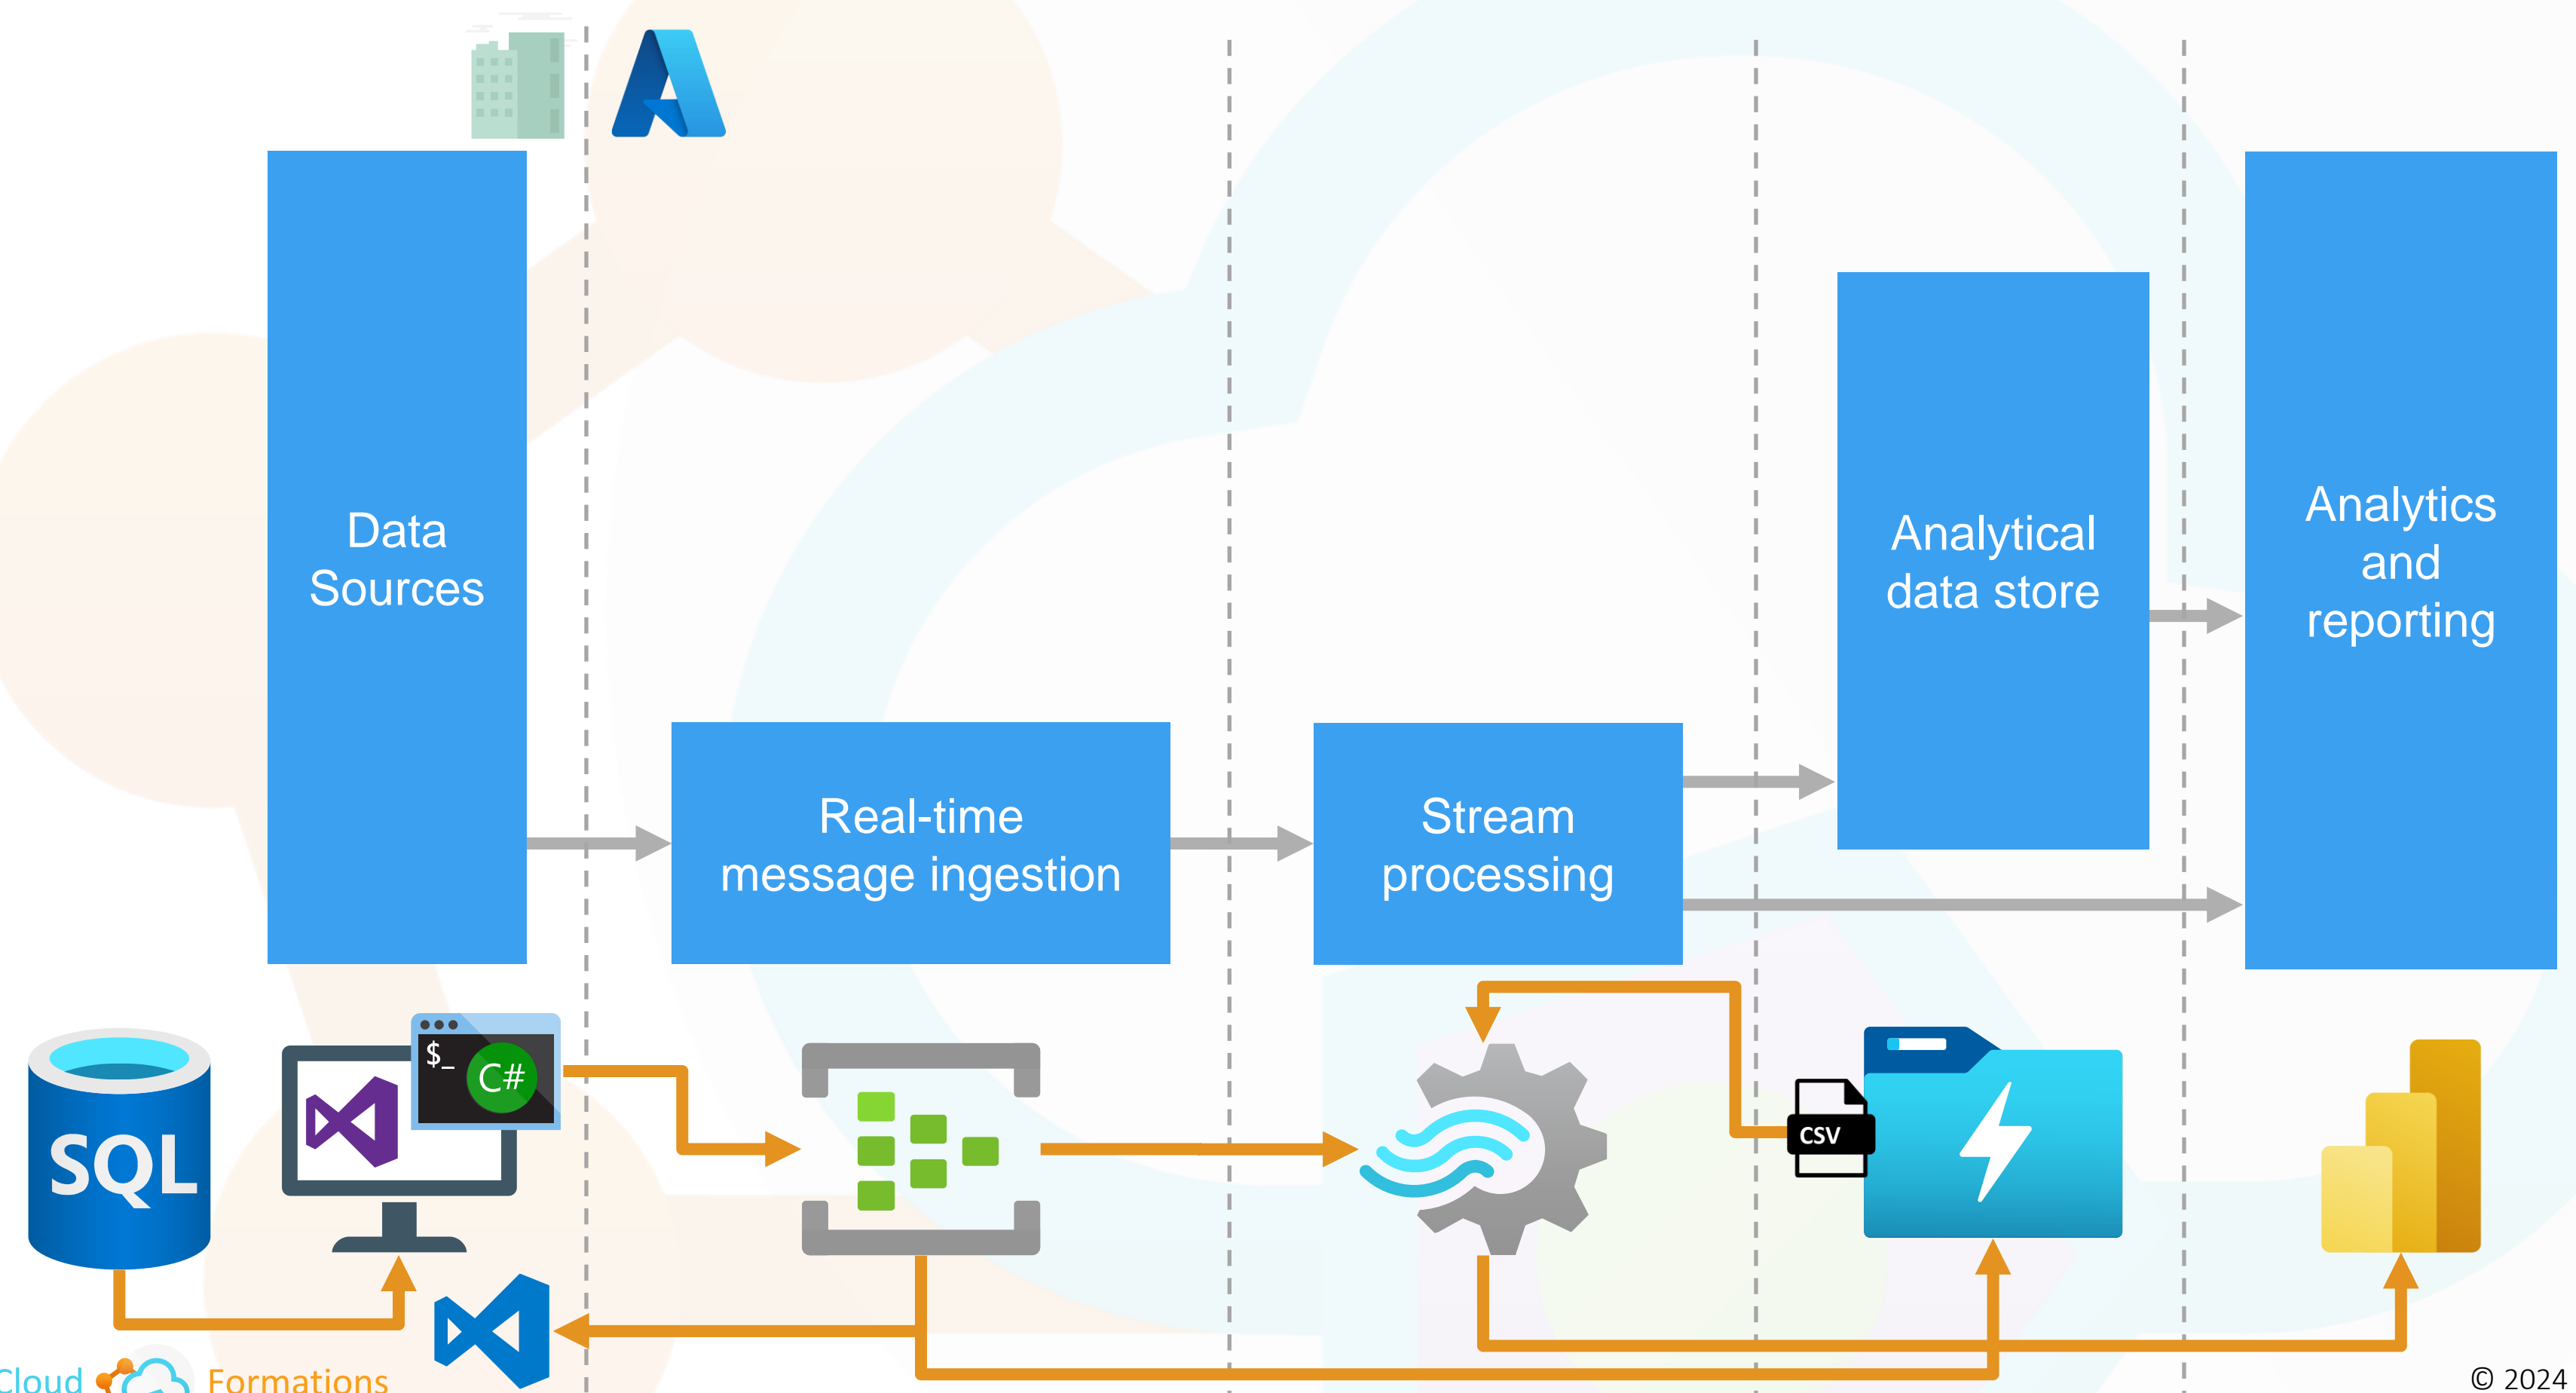
Fabric Real-Time Data Handling




Fabric Real-Time Data Handling (Mirroring)



Azure Real-Time Data Handling




Fabric Real-Time Data Handling



Conclusions

Cloud Formations



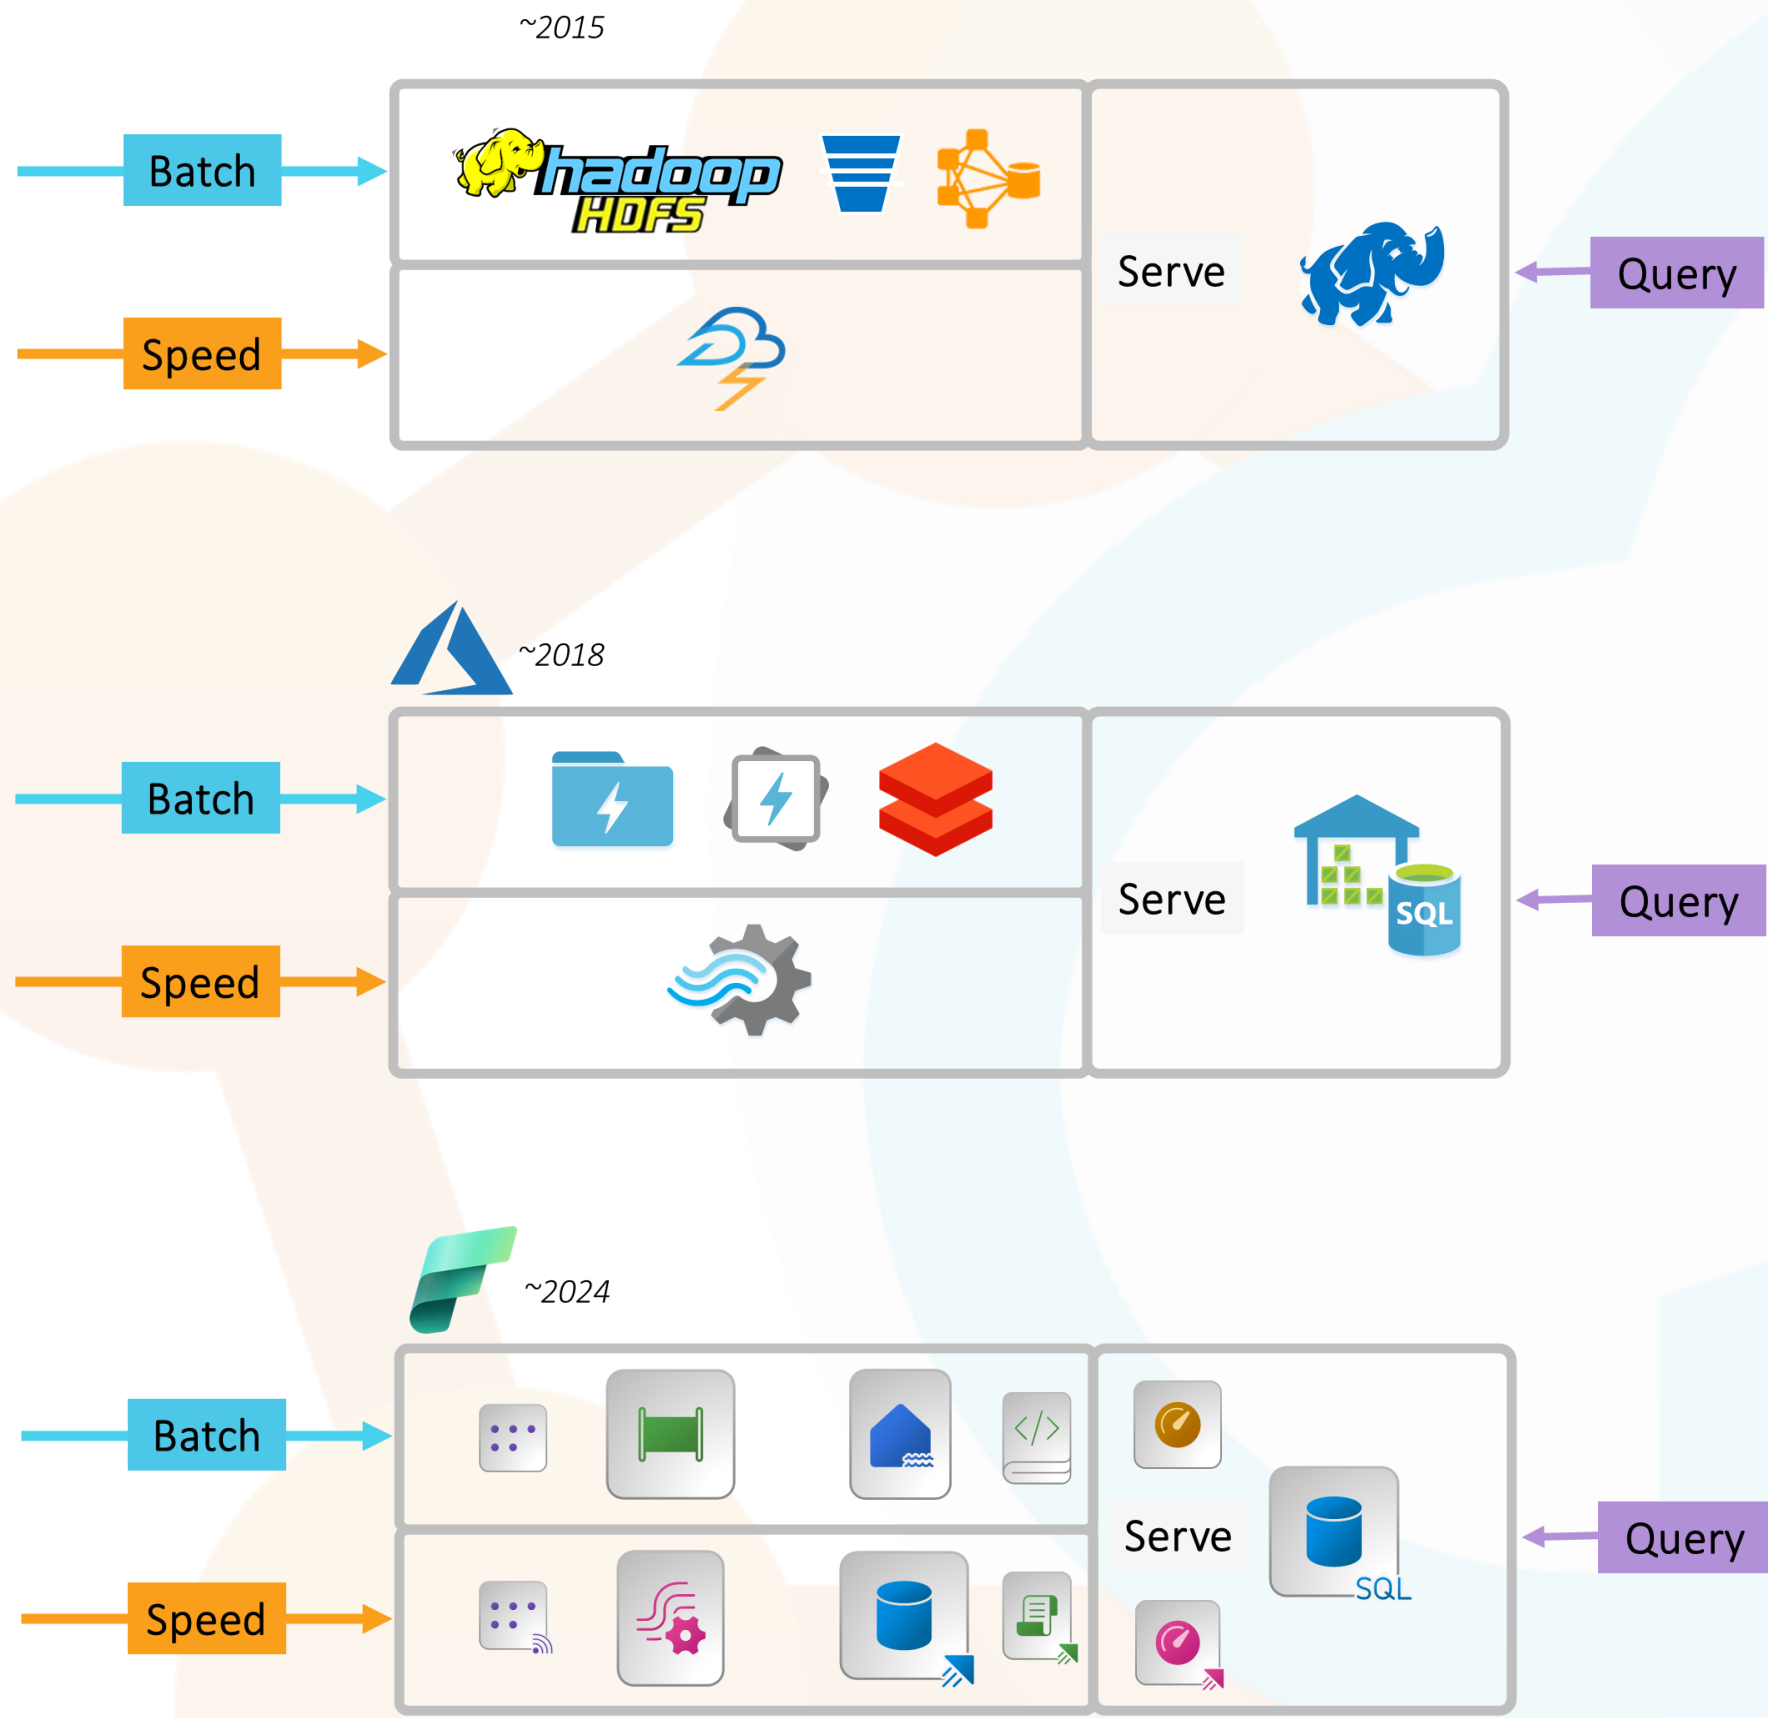
Glossary

Term	Definition
Big Data	Any data that you cannot process in the time that you have/want using the technology you have.
Real-time Data	Delivering data from the producer to consumer as fast as possible using the technology you have.
Near Real-time Data	Delivering data from the producer to consumer within 1 minute of it being created.
Data Stream	Data that is constantly flowing from producer to consumer in near real-time.

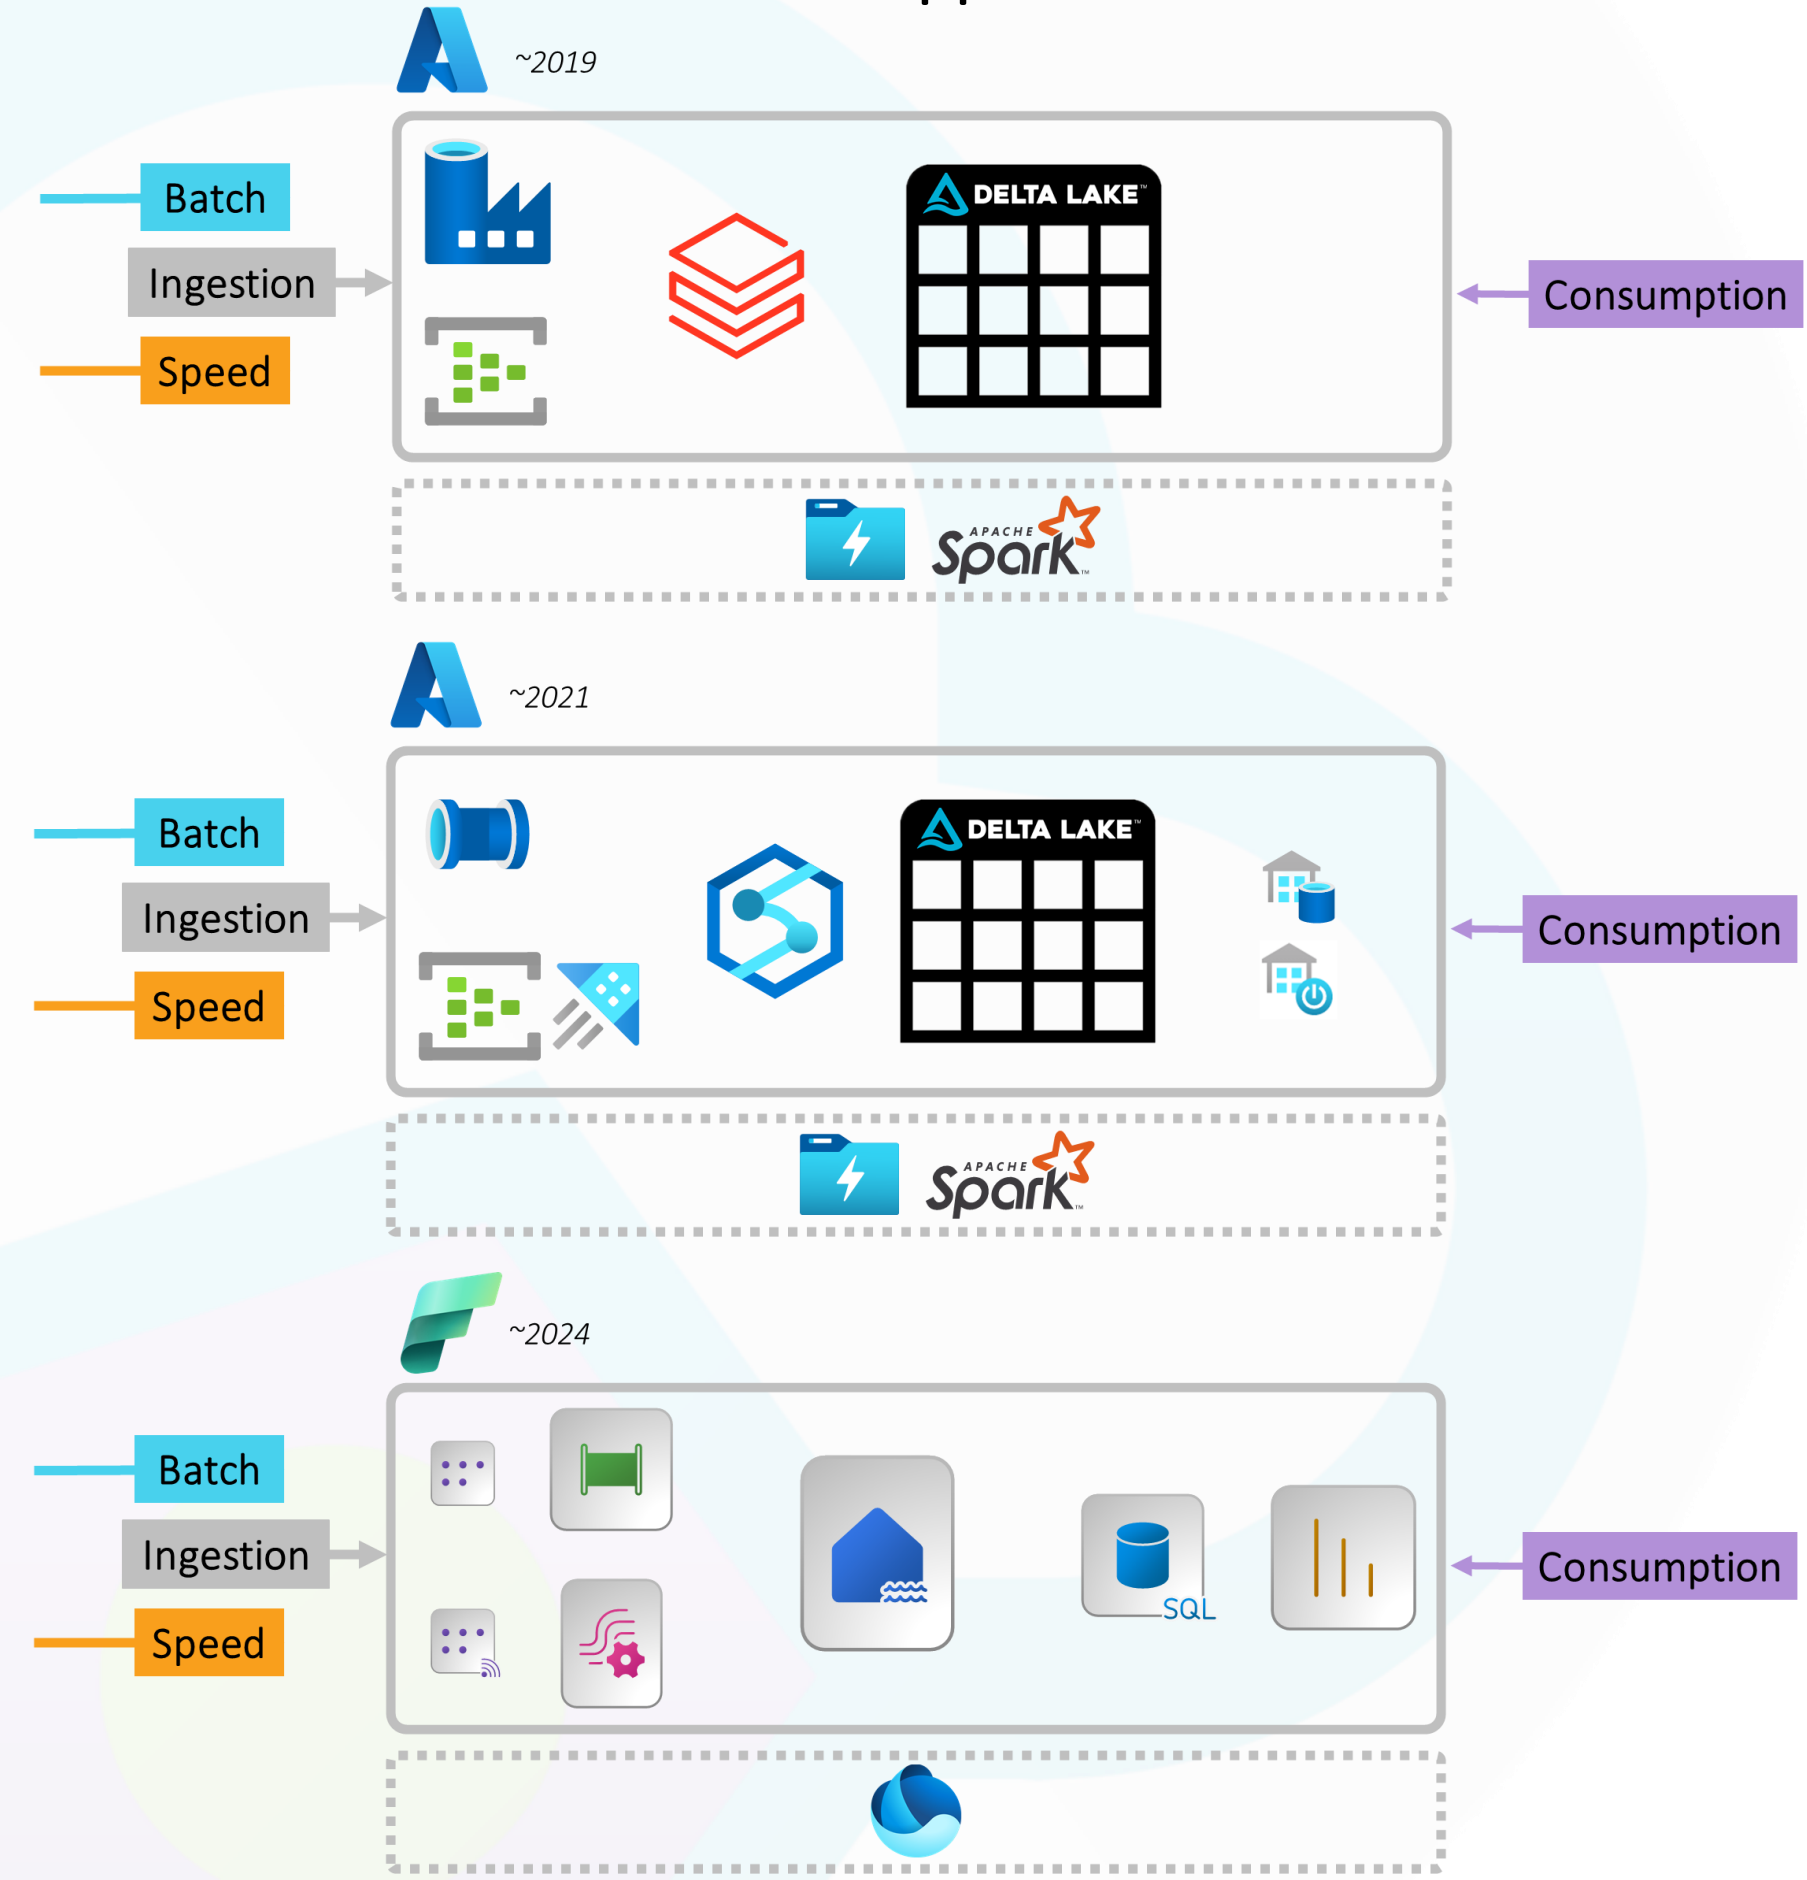
Lambda & Kappa Architectures vs Technology



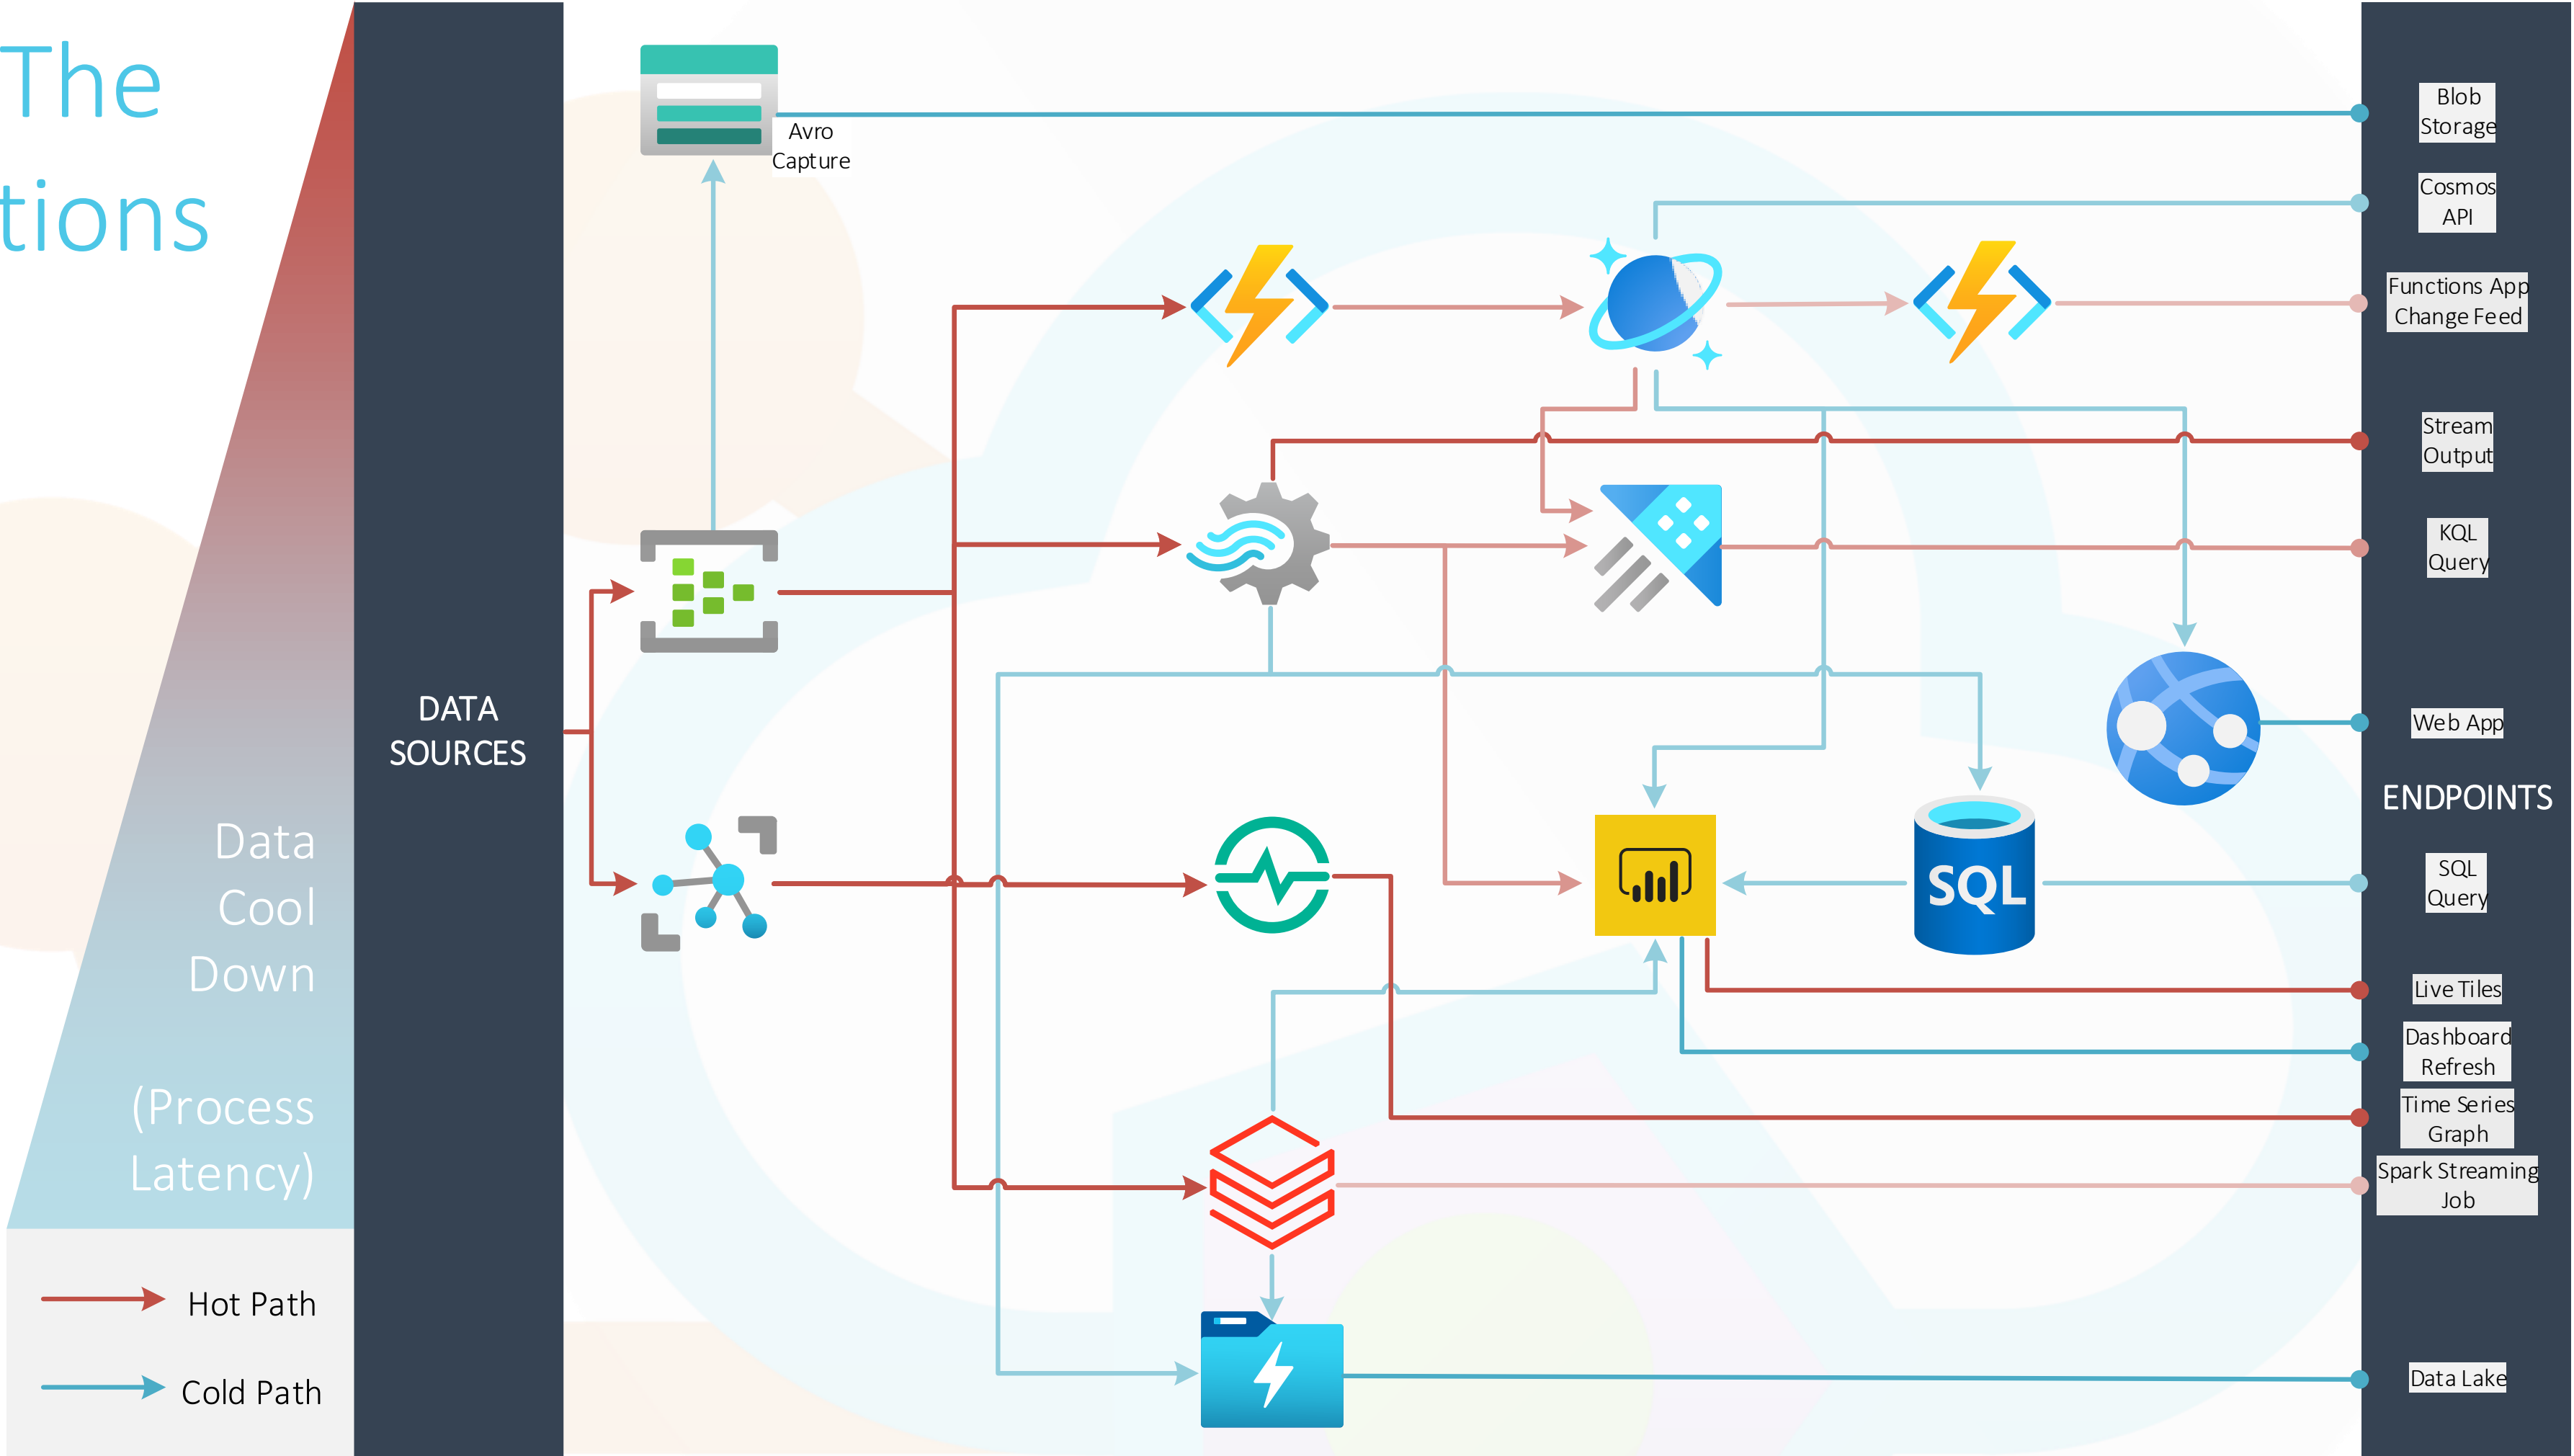
Lambda



Kappa



All The Options



Customer Case Studies & Testimonies

Implementing a Metadata-Driven Data Lakehouse on Azure with CF.Cumulus

Leveraging ready-to-go pipelines, connections and metadata to accelerate the delivery process of a cloud data platform.

Cloud Formations utilised our open-source metadata-driven Lakehouse data platform accelerator, CF.Cumulus, to rapidly ingest data from Oracle and Dynamics 365, leveraging Microsoft Analytics Platform technologies to deliver use cases for a large London based housing association.

Technologies
Using Microsoft Azure cloud native resources, a data mesh architecture is being designed using the following compute and storage components:

- Data Lake
- Data Factory
- Databricks
- SQL Database
- Functions
- Microsoft Fabric
- Azure DevOps

Solution Overview
Using the CF.Cumulus accelerator provided rapid deployment of a best practice Lakehouse data platform with seamless integration, meeting networking and governance requirements of the client. The pre-built product connections and out-of-the-box CF.Cumulus pipelines in Azure Data Factory facilitated deeper collaboration with the client and enabling self-service of the platform.

Data is then served through Microsoft Fabric for a unified business user experience meeting the strategic ambitions of the organisation and enable consumers to explore and self-serve from trusted, curated datasets. The solution uses the best of all products to deliver a complete 'medallion architecture'.

Key Features

- Metadata-driven data ingestion framework.
- Intelligent data loading through CDC logic and data load history.
- Cost-efficient compute implementation for both development and production experiences.
- Resilient and restartable orchestration pipelines.

Outcome
The solution provided the business with:

- Datasets ingested from source with little configuration efforts by development team.
- Repeatable framework and training provided to ingest additional data sources.

Next Steps
Reach out to learn more about how we could implement this architecture in your organisation via our contact details (top right) or please book some time in our calendar via the following link: bit.ly/cf-chat

Our Services

- Data Platform Design
- CF.Cumulus Deployment
- Technical Oversight
- Architecture Review
- End to End Training

<https://bit.ly/cf-chat>

Delivering Near Real-Time Satellite Internet Analytics Using Microsoft Fabric Paired With Microsoft Azure for Scaled Data Ingestion

Cloud Formations teamed up with Onyx Data to lead the technical design and delivery of a near real-time analytics solution, providing satellite internet (Starlink) usage telemetry to a global service provider.

Solution Overview
The greenfield implementation follows a microservices approach to data handling, leveraging **Azure Functions Apps** to ingest telemetry from the Starlink API at scale and metadata driven, feeding into a **Microsoft Fabric Event Stream**. Bootstrapped by **Integration Pipelines**, data feeds into Realtime Analytics Kusto Databases and Lakehouse structures for reporting using **Power BI** Dashboards.

The platform was managed using a combination of **Azure DevOps** release pipelines, **Git** source control and **Agile** backlogs. Chained together with Microsoft Fabric artifacts deployed using **Workspace Pipelines**.

Use Cases
The solution dashboards keep internet service users informed on:

- Available bandwidth.
- Aggregated data transfers (upload/download).
- Provider usage limits.
- Connection latency and obstructions.

Time Scales
With an expert team led by Microsoft Data Platform MVPs, this solution went from design to delivery in just **2 months**.

Our Services

- Data Platform Design/Build
- Technical Oversight
- Architecture Review
- Data & AI Strategy
- End to End Training

<https://bit.ly/cf-chat>

Delivering 'The Golden Thread' Dataset Using Data Mesh Concepts Supporting Housing & Community Development

Cloud Formations have partnered with a **G15 Housing Association** to pioneer a solution that not only embraces the 'Golden Thread' use case but also revolutionises the Microsoft Data Platform landscape.

What is the Golden Thread? Building Regulations Advisory Committee: golden thread report. See full article: bit.ly/ref-article

The Challenge
The Housing Association industry's data landscape often resembles a labyrinth, hindering insights and innovation. Our client, a leading housing association, faced a critical challenge – how to harness their data unlocking the potential for transformative outcomes while navigating the complexities of the 'Golden Thread' use case. Amongst other business critical deliverables.

Data Mesh Principles
The principles of a data mesh are well defined within the technical community, but often left open to interpretation for businesses seeking to apply them.

- Domain-oriented decentralised data ownership and architecture.
- Data as a product.
- Self-serve data infrastructure as a platform
- Federated computational governance.

The Technology Stack
Using Microsoft Azure cloud native resources, a data mesh architecture is being designed using the following compute and storage components:

- Data Lake
- Data Factory
- Databricks
- SQL Database
- Purview
- Functions
- Microsoft Fabric
- Power BI

All connect and leading to the use of governance as an enabler to do more with data safely.

See Also
Our Data Mesh Reference Architecture, plus others (inspired by this project work). bit.ly/cf-refarch

Outcome
Design a future proof, decentralised data mesh architecture to serve all data requirements across the business, including support for the 'Golden Thread' use case. A platform that will empower business users, unlock value in data assets, federate data governance, feed the overall business data strategy and re-organise siloed data products.

Next Steps
Reach out to learn more about how we could implement this architecture in your organisation via the contact details (top right) or please book some time in our calendar via the following link: bit.ly/cf-bookings

Our Services

- Data Strategy
- Architecture Review & Oversight
- Data Culture & Adoption
- AI Strategy & Consulting
- Data Mesh & Fabric Adoption
- Data Science as a Service
- Training

<https://bit.ly/cf-bookings>

Delivering Real-Time Marketing Analytics Using Microsoft Fabric Outreach Intelligence & Campaign Interactions

Cloud Formations partnered with a management consultancy to deliver a complete set of marketing analytics dashboards based on ongoing outreach initiatives, using industry standard data sources.

Solution Overview
Near real-time streaming of marketing outreach data into Power BI dashboards to provide up to the minute analytics on the performance of campaigns to inform future content and audience enrolment.

The Technology Stack
Using Microsoft Azure cloud native resources, a data mesh architecture is being designed using the following compute and storage components:

- Fabric – Integration
- Fabric – Lakehouse
- Fabric – Event Stream
- Fabric – Notebooks
- Functions App
- Storage Account
- Key Vault
- Power BI

Lambda Architecture Principles
The solution leverages Microsoft Fabric to transform and consolidate high-frequency streaming data with a batch feed into a curated layer ready for reporting.

See Also
Our Data Mesh Reference Architecture, plus others (inspired by this project work). bit.ly/cf-refarch

Outcome
Live operational dashboards to support the Sales Development team with up-to-the-minute data on performance. This allows the team to rapidly test, adjust and optimise sales outreach activity over a time horizon measured in minutes rather than days, massively improving the team's productivity and results.

Next Steps
Reach out to learn more about how we could implement this architecture in your organisation via our contact details (top right) or please book some time in our calendar via the following link: bit.ly/cf-chat

Our Services

- Data Strategy
- Architecture Review & Oversight
- Data Culture & Adoption
- AI Strategy & Consulting
- Data Mesh & Fabric Adoption
- Data Science as a Service
- Training

<https://bit.ly/cf-chat>

Automated Mailbox Attachment File Ingestion & Processing Using Microsoft Fabric Low Code Data Flows

Cloud Formations assisted a leading utilities provider in creating a low code data ingestion process from an Office 365 mailbox, with built-in data quality checks to minimise the human effort of investigating the quality of files sent to the inbox. Surfacing results in a Microsoft Fabric Lakehouse Workspace.

Solution Overview
This implementation required tooling familiar to the Finance Team which have ownership of the mailbox contents, and the process. As such, Power Query was used via **Microsoft Fabric Dataflows** to create a low code process for extracting the relevant files from the mailbox in a reproducible and interpretable way. **Microsoft Fabric Data Pipelines** were then used to bootstrap the activities and provide the basis for the required data quality checks and subsequent alerting. Once ingested all data was served through a SQL Analytics Endpoint to users via a **Microsoft Fabric Lakehouse** view offering data to consumers of the compliant files in a Silver data layer ready for further modelling and aggregation is Power BI.

Outcome
The solution framework allowed the team to:

- Provide business users with a unified view of validated data received in the inbox.
- Mitigate manual quality troubleshooting tasks.
- Work with familiar tools with little upskilling requirements.

Time Scales
The implementation was produced within a single technical sprint of 10 days, focussed on the following stages:

Next Steps
Reach out to learn more about how we could implement this architecture in your organisation via our contact details (top right) or please book some time in our calendar via the following link: bit.ly/cf-chat

Our Services

- Data Platform Design/Build
- Technical Oversight
- Architecture Review
- Data & AI Strategy
- End to End Training

<https://bit.ly/cf-chat>

“Cloud Formations provide insight and pragmatic advice based on deep understanding and experience across a range of technical domains. I value their sleeves rolled up approach and willingness to flex across a range of requirements.”

“Cloud Formations recently facilitated a value discovery session with us using their engagement framework. The insights and outputs from this session have proven incredibly valuable and we are excited to collaborate with them on the next steps in our journey.”

“We have partnered with Cloud Formations and its team of SMEs for several years on many different projects. Having them as a trusted advisor to inform data and technical strategy has been incredibly valuable. They have and continue to support to us in the design and delivery of a data mesh platform allowing us to scale our data products for affective business user consumption.”




The screenshot shows the Cloud Formations website with a navigation menu (Home, About, Services, Resources, Careers, Contact Us) and a main heading 'Community Content'. Below the heading is a large image of a conference audience. A specific content item is highlighted with a grey background, containing the following text:


- How Can Microsoft Fabric Have an Impact on Your Business?**
- Speaker:** Paul Andrew
- Date:** 13 July 2023
- Hosts:** Microsoft Fabric- UK User Group

Below the text is a thumbnail image of a presentation slide with the same title and speaker information. A blue button labeled 'PDF of Slides' is located at the bottom left of the content item. An orange arrow points from this button to a PDF icon on the right side of the slide.



<https://www.cloudformations.org/community-content>





Thank You

- ✉ paul@mrpaulandrew.com
- [in](https://www.linkedin.com/in/mrpaulandrew) [In/mrpaulandrew](https://www.linkedin.com/in/mrpaulandrew)
- [@mrpaulandrew](https://twitter.com/mrpaulandrew)

Paul Andrew
Technical Architect



Cloud Formations